

THE
INHABITED
BODY

When the Partnership Breaks



Dr Horst Herb

BACTERIA · ARCHAEA · EUKARYA

The Inhabited Body — Volume 2: Feeding, Fixing, and Fighting Your Microbiome

What We Do to Our Microbiome — and What It Does Back

Horst Herb

2026-04-22

Contents

The Inhabited Body — Volume 2: Feeding, Fixing, and Fighting Your Microbiome 6

Chapter 18: Antibiotics — Collateral Damage 7

18.1 The Broad-Spectrum Problem: Carpet-Bombing an Ecosystem	7
18.2 Short-Term Disruption and Long-Term Scars	9
18.3 Early-Life Antibiotic Exposure and Lifetime Consequences	11
18.4 Antibiotic Resistance Genes in the Commensal Microbiome — A Hidden Reservoir .	13
18.5 <i>Candida</i> Blooms and the Kingdom That Fills the Vacuum	15
18.6 The Recovery Phase — What Actually Happens When Antibiotics Stop	16
18.7 Towards Precision: Narrow-Spectrum and Microbiome-Sparing Approaches	18

Chapter 19: Phage Therapy — The Comeback 23

19.1 A History: D'Hérelle, Soviet Medicine, and the Road Not Taken	23
19.2 Two Patients, Two Infections, Two Worlds	25
19.3 Landmark Cases: When Compassionate Use Changed the Conversation	28
19.4 The Limits of the Comeback	30

Chapter 20: Diet — Feeding the Ecosystem 33

20.1 The Dry Season in Hadzaland	33
20.2 The Gardener's Dilemma	33
20.3 The Currency of the Realm: Fibre and the Microbiota-Accessible Carbohydrates . . .	35
20.4 Diversity Beats Dose: The Thirty-Plant Rule	39
20.5 The Mediterranean Pattern, Put to the Test	43
20.6 Fermented Foods: The Other Living Diet	46
20.7 Polyphenols: Why It Depends Whose Garden You Have	49
20.8 The Other Side of the Plate: Protein, Fat, and the Metabolites You Don't Want . . .	51
20.9 Ultra-Processed Food: When the Package Is the Problem	54
20.10 When You Eat: The Circadian Microbiome	57
20.11 Personalised Nutrition: Why "What Should I Eat?" Is the Wrong Question	59
20.12 The Controversies, Met Head-On	62
20.13 What This Means for Tuesday Dinner	66
20.14 Looking Ahead	68

Key References	68
Chapter 21: The Other Levers — How the Rest of How We Live Shapes the Ecosystem Inside Us	73
21.1 The Garden Has Weather	73
21.2 The Forgotten Clock: Artificial Light and the Diurnal Microbiome	74
21.3 Twenty-Two Degrees All Year: The Thermostat and the Unfelt Season	76
21.4 The Invisible Garden in the Air: HEPA Filters, UVC Lamps, and the Post-Pandemic Reflex	78
21.5 Dirt, Biodiversity, and the Old Friends We Stopped Meeting	80
21.6 The Dog in the Doorway: Household Pets and the Infant Microbiome	83
21.7 The Household as a Microbial Unit: What We Share With the People We Live With	84
21.8 Moving the Body: Exercise as a Microbial Intervention	86
21.9 Sleep and Stress: The Nervous System’s Dialogue With the Gut	87
21.10 The Medicine Cabinet: Non-Antibiotic Drugs and the Microbiome	88
21.11 The Water We Drink: Chlorine, Hardness, and the Microbes That Aren’t There	90
21.12 Taking Stock: The Modern Environment as a Microbiome Experiment	91
Key References	93
Chapter 22: Drugs Beyond Antibiotics — The Quiet Pharmacology of the Gut	96
22.1 The Quietest Side Effect	96
22.2 Reading the Evidence: A Note Before We Start	97
22.3 Closing the Acid Gate: PPIs and Their Cousins	98
22.4 The Drug That Works Through Microbes: Metformin	101
22.5 The Quiet Erosion: NSAIDs and the Small Bowel You Cannot Feel	103
22.6 Statins, Cholesterol, and the Bact2 Enterotype	106
22.7 The Mind-Altering Drugs: Antipsychotics, Antidepressants, and the Microbial Cost of Calm	108
22.8 Opioids and the Bowel That Forgets to Move	111
22.9 The Forgotten Effect of Laxatives	112
22.10 Hormones in a Pill: Contraceptives, HRT, and the Estrobolome	113
22.11 The Supplement Aisle: Iron, Vitamin D, Fish Oil, and What’s Real About Berberine	115
22.12 Polypharmacy: The Cumulative Tax	118
22.13 Pharmacomicrobiomics: When Your Microbes Decide What Your Medicine Does	120
22.14 What This Means for the Person Reading This Book	122
References for Chapter 22	123
Chapter 23: Probiotics — Promise, Hype, and Hard Truths	127
23.1 What “Probiotic” Actually Means — and the Gap Between Definition and Marketplace	127
23.2 The Evidence Hierarchy: Which Strains, for Which Conditions, Actually Work?	129
23.3 The Colonisation Problem — Most Probiotics Don’t Establish Residence	132
23.4 The Personalisation Problem — Mucosal Colonisation Varies Dramatically Between Individuals	133
23.5 When Probiotics Backfire: The Post-Antibiotic Paradox	134
23.6 The Competition Framework: “Beneficial” Species as Ecological Disruptors	136
23.7 The Infant Probiotic Question — NEC Prevention vs Ecological Disruption	138
23.8 Probiotics and the Resistome: Can Probiotics Expand Antibiotic Resistance?	139
23.9 Next-Generation Probiotics: <i>Akkermansia</i> , <i>Faecalibacterium</i> , Engineered Strains	141

23.10 Regulatory Grey Zones, the Supplement Industry, and What the Consumer Should Know	144
23.11 A Measured Conclusion: Not “Probiotics Are Bad” but “Probiotics Are Drugs We’re Prescribing Without a Diagnosis”	145
Key Chapter References	147
Chapter 24: Faecal Microbiota Transplantation — The Nuclear Option	149
24.1 A Title That Has to Be Earned	149
24.2 Four Centuries of Yellow Soup, and What We Can Honestly Claim About It	150
24.3 Amsterdam, 2013: The Trial That Stopped Itself	151
24.4 What Actually Happens: Inside a Modern FMT Clinic	153
24.5 Beyond <i>C. difficile</i> : The Frontier That Does Not Quite Work Yet	156
24.6 The Concept of the Super-Donor	159
24.7 The Safety Signal We Have to Take Seriously	160
24.8 The Underground: DIY-FMT and What It Tells Us	161
24.9 From Stool to Standard Product: The End of FMT As We Know It?	164
Key Chapter References	166
Chapter 25: Precision Microbiome Engineering — From Nuclear Option to Scalpel	168
25.1 The Opposite of a Transplant	168
25.2 Four Kinds of Precision	169
25.3 What Already Works: The Defined Consortia	170
25.4 The Engineered Commensals: Synlogic and the Bacteria That Eat Poisons	171
25.5 The CRISPR Antimicrobials: Killing One Species at a Time	173
25.6 The Regulatory Landscape: Why Approval Is Slower Than the Science	175
25.7 The Risks Specific to Precision Engineering	177
25.8 What Is Plausibly Near, and What Is Not	179
25.9 Scalpel, Not Nuclear: A Final Framing	180
Key Chapter References	181
Chapter 26: Inflammatory and Autoimmune Disease	182
26.0 The Twentieth-Century Curve	182
26.1 Inflammatory Bowel Disease: The Paradigm Condition	183
26.2 The Fungal Dimension: The ASCA Mystery, Thirty Years On	187
26.3 Type 1 Diabetes and the “Leaky Gut” Hypothesis	190
26.4 Multiple Sclerosis and Rheumatoid Arthritis: Autoimmunity at a Distance	193
26.5 Coeliac Disease: Genes, Gluten, and Microbes	197
26.6 A Common Mechanism or a Convergent Phenotype?	200
Key Chapter References	203
Chapter 27: Metabolic Disease and Obesity	206
27.0 The Weight of the Ecosystem	206
27.1 The Famous “Obese Microbiome” Experiments — What They Showed and What They Didn’t	207
27.2 Energy Harvest, Bile Acids, and Appetite Signalling	209
27.3 Type 2 Diabetes, Fatty Liver Disease, and Cardiovascular Risk	212
27.4 <i>Akkermansia muciniphila</i> and the Next-Generation Probiotic Hope	216
27.5 Why Microbiome-Based Weight Loss Therapies Haven’t Arrived Yet	217

Key Chapter References	219
Chapter 28: Cancer	221
28.0 The Oldest Suspicion	221
28.1 <i>Helicobacter pylori</i> and Gastric Cancer — The Original Microbe-Cancer Link	222
28.2 <i>Fusobacterium nucleatum</i> and Colorectal Cancer	224
28.3 The Microbiome as a Modulator of Chemotherapy	227
28.4 Immune Checkpoint Inhibitors: The Microbiome Predicts Response	228
28.5 Predicting Cancer Recurrence: The Microbial Fingerprint	230
28.6 Could We Use the Microbiome for Early Cancer Detection?	231
A Relationship Still Unfolding	232
Chapter 29: Neurological and Psychiatric Disease	236
29.0 The Hardest Organ	236
29.1 Parkinson's Disease: Does It Start in the Gut?	237
29.2 Alzheimer's Disease and Neuroinflammation	242
29.3 Depression and the Promise of Microbiome-Targeted Interventions	246
29.4 The Translation Gap: From Mouse Models to Human Trials	255
29.5 Closing: A Field in Adolescence	260
References	260
Chapter 30: The Microbiome as a Diagnostic Tool	265
30.0 Reading the Ecosystem	265
30.1 Microbial Biomarkers for Disease Screening	266
30.2 Companion Diagnostics: Predicting Drug Response	273
30.3 Sensors, Capsules, and Real-Time Monitoring	274
30.4 AI and the Microbiome	277
30.5 The Direct-to-Consumer Problem	280
30.6 What Is Ready, What Is Not, and What Is Coming	282
Key References	283
The Inhabited Body — Bibliography by Chapter	288
Ch. 1 — You Are Never Alone	288
Ch. 2 — Built Together: How Microbes Shaped Human Evolution	288
Ch. 3 — Your First Inheritance: Colonisation Across a Lifetime	289
Ch. 4 — Tools of Discovery	289
Ch. 5 — The Gut	289
Ch. 6 — The Skin	291
Ch. 7 — The Respiratory Tract	291
Ch. 8 — The Urogenital Microbiome	293
Ch. 9 — The Oral Microbiome	294
Ch. 10 — The Mycobiome	294
Ch. 11 — The Brain Microbiome	295
Ch. 13 — The Gut-Brain Axis	295
Ch. 14 — Behavioural Modulation	296
Ch. 15 — Metabolism, Hormones, and the Endocrine Microbiome	297
Ch. 16–17 — Bacteriophages	297
Ch. 18 — Phage Therapy	297

Ch. 19 — Antibiotics: Collateral Damage	297
Ch. 20 — Diet: Feeding the Ecosystem	298
Ch. 21 — Environment, Lifestyle, and the Modern World	301
Ch. 22 — Drugs Beyond Antibiotics	301
Ch. 23 — Probiotics: Promise, Hype, and Hard Truths	305
Ch. 24 — Faecal Microbiota Transplantation	305
Ch. 25 — Precision Microbiome Engineering	306
Ch. 26 — Inflammatory and Autoimmune Disease	307
Ch. 27 — Metabolic Disease and Obesity	309
Ch. 28 — Cancer	311
Ch. 29 — Neurological and Psychiatric Disease	312
Ch. 30 — The Microbiome as a Diagnostic Tool	312
Ch. 31 — Ethical, Social, and Philosophical Questions	313
Notable Observations	314
Ch. 13 — The Gut-Brain Axis: Microbes and the Mind	314
Ch. 17 — Phages, Immunity, and the Tripartite Relationship	315

The Inhabited Body — Volume 2: Feeding, Fixing, and Fighting Your Microbiome

What We Do to Our Microbiome — and What It Does Back

Horst Herb

☐ **PROOF COPY — 2026-04-22 — NOT FOR DISTRIBUTION**

Chapter 18: Antibiotics — Collateral Damage

18.1 The Broad-Spectrum Problem: Carpet-Bombing an Ecosystem

On the evening of 11 December 1945, Alexander Fleming stood at a lectern in Stockholm to deliver his Nobel Lecture. He had already said what the audience came to hear — the story of the contaminated Petri dish, the ring of inhibition around the *Penicillium* colony, the miracle drug that had just finished saving Allied soldiers from wound infections in Europe and the Pacific. But he ended his lecture with a warning that the audience, exhausted by war and intoxicated by the promise of antibiotics, was in no mood to take seriously.

“It is not difficult,” Fleming said, “to make microbes resistant to penicillin in the laboratory by exposing them to concentrations not sufficient to kill them, and the same thing has occasionally happened in the body.” He went on to describe, with a kind of clinical premonition, how a man who bought penicillin at the chemist’s and took it for a sore throat might give his streptococci enough exposure to select for resistant mutants, then pass those mutants to his wife, who might die of a resistant pneumonia — “and who is responsible for that lady’s death? Why, the man who negligently played with penicillin” [REF:fleming1945].

It was the first public warning about antibiotic resistance, delivered by the discoverer himself, in the year of penicillin’s greatest triumph. It was ignored. For the next seventy years, antibiotics were prescribed as if the supply were infinite, the pathogens static, and the ecosystem they passed through empty.

That ecosystem, of course, was not empty. By the time a course of amoxicillin reaches the small bowel, it enters a community of hundreds of bacterial species, living in structured relationships with the mucus layer, with each other, and with the host. The drug does not distinguish between the pathogen the physician was worrying about and the commensals the physician never thought about. It kills what it can reach, in the concentrations it can reach, and it leaves the survivors to fill the vacancies. What happens next is the subject of this chapter.

The Ecologist’s Objection

If an ecologist were asked to describe modern antibiotic therapy in ecological terms, they would probably reach for a metaphor from forestry. Imagine a mature, old-growth forest — centuries in the making, with thousands of species arranged in layered niches, each holding the others in check through competition, predation, and mutualism. Now imagine treating that forest for a bark-beetle outbreak by aerial spraying with a broad-spectrum herbicide. You would certainly kill the bark beetles’ host trees, and thereby the beetles, but you would also kill most of the other trees, most of the understory, most of the fungi, most of the insects, and — indirectly — most of the animals that depend on any of these. The bark-beetle problem would be solved. The forest would be gone.

This is not a strained analogy. When the oncologist Martin Blaser, working at New York University, began asking in the mid-2000s whether the antibiotic era had been quietly producing an ecological catastrophe inside human bodies, the forest metaphor was one he returned to repeatedly [REF:blaser2014]. His argument, developed over a decade of research and eventually published as the book *Missing Microbes*, was that broad-spectrum antibiotics do not merely treat an infection —

they depopulate a habitat, and the recolonisation that follows is not always complete. Some species come back. Some do not. The forest that regrows is not the forest that was there before.

The empirical question behind Blaser's hypothesis is simple: when you finish a course of antibiotics, does your microbiome really return to baseline? For decades, the assumption was yes. Bowel flora were thought to be resilient, self-restoring, and more or less interchangeable between individuals. The few studies that looked at the question used culture methods that could see only a fraction of the community, and they generally reported that things looked fine after a week or two. The molecular methods we described in Chapter 4 changed that picture completely.

What a Single Course Actually Does

The landmark study on this question was published in 2018 by Albert Palleja and colleagues at the University of Copenhagen. The design was elegant and unsettling. Twelve healthy young men — average age twenty-four, no recent antibiotic use, no underlying illness — were given a four-day course of three broad-spectrum antibiotics: meropenem, gentamicin, and vancomycin. This combination was chosen to mimic the regimen a patient might receive for a serious hospital-acquired infection. The volunteers provided stool samples at baseline, at the end of treatment, and then at eight time points over the following six months. The samples were analysed by deep shotgun metagenomic sequencing, the technique that allows researchers to identify every microbial species present and estimate its abundance without the biases of culture [REF:palleja2018].

The headline result, repeated in news coverage at the time, was reassuring: most of the bacterial species present at baseline had returned to detectable levels by day forty-two. The microbiome was “resilient.” The word *resilient* did a lot of work in that coverage. What it obscured was everything in the paper that came after the headline.

First, “most” was not “all.” Nine common species that had been present in all twelve subjects at baseline were still undetectable in most subjects six months later. These were not obscure organisms. They included species of *Bifidobacterium adolescentis*, *Coprococcus eutactus*, and *Methanobrevibacter smithii* — species associated with short-chain fatty acid production, immune regulation, and hydrogen metabolism in the colon. The researchers' wording was cautious: these species had not recovered “during the observation period.” The implication, unstated but clear, was that they might not recover at all.

Second, the community that did come back was not the community that had left. Overall diversity — the number of species present, weighted by their relative abundances — remained below baseline at six months. More importantly, the *functional* composition of the community had shifted. Genes involved in virulence, in biofilm formation, and in antibiotic resistance were enriched in post-treatment samples. The bacteria that had survived the antibiotic onslaught and expanded into the vacated niches were, on average, more aggressive and less cooperative than the ones they had replaced. The forest had regrown, but it was a forest of weeds.

Third — and this was the finding that made microbiome researchers most uncomfortable — the post-treatment communities contained forty-seven new genes conferring resistance to a variety of antibiotics, including some that the volunteers had never received. These were not novel mutations. They were resistance genes that had been present all along, at low frequency, in minority members of the community. The antibiotic treatment had acted as a selective sieve, enriching for the carriers. When the community regrew, it regrew from a founder population that was already more antibiotic-resistant than the one it replaced.

One four-day course. Twelve healthy young men. A six-month shadow.

The Concentration Gradient That Nobody Prescribes

To understand why antibiotics do this much collateral damage, it helps to think about where the drug actually goes when you swallow it. When a patient takes an oral antibiotic, the drug is absorbed across the small intestinal wall into the bloodstream, from which it is distributed to the site of infection — a lung, a kidney, a surgical wound. But absorption is never complete, and excretion is never instantaneous. A substantial fraction of the dose passes through the small intestine unabsorbed and arrives in the large intestine, where the densest part of the microbiome lives. A further fraction, already absorbed, is secreted back into the gut via the bile — the so-called enterohepatic circulation. For some antibiotics, the concentration that reaches the colon is a small fraction of the blood level; for others, it is much higher than anywhere else in the body.

This means that the dense bacterial populations of the colon — *Bacteroides*, *Faecalibacterium*, *Roseburia*, *Bifidobacterium*, and the hundreds of less-abundant species that live alongside them — are exposed to antibiotic concentrations they did not need, in pursuit of a pathogen that is almost never in the colon. The “target” of the prescription is a streptococcus in the throat, a *Staphylococcus aureus* in a wound, an *E. coli* in the bladder. The *victims* of the prescription are the trillions of bystanders who happen to share the same body.

Pharmacologists have a phrase for this kind of exposure: *off-target effect*. It is the same phrase used when a cancer drug damages healthy tissue, or when a psychiatric drug disturbs sleep. But there is an important difference. When a chemotherapy drug damages the bone marrow, clinicians monitor the damage and stop if it becomes too severe. When an antibiotic damages the gut microbiome, nobody is looking. The prescription is written for the infection. The bystander community is neither monitored before nor counted after. In the vocabulary of clinical medicine, the microbiome is not yet an organ; it is collateral.

18.2 Short-Term Disruption and Long-Term Scars

The phrase “collateral damage” suggests something transient, the way a bruise is transient. And in the short term, that is how antibiotic disruption feels to most patients — a few days of loose stools, perhaps a bit of bloating, occasionally a yeast infection, and then back to normal. The “back to normal” part is an illusion, but it is a convincing one, because the host-level symptoms usually resolve even when the community underneath has not.

To see what the damage actually looks like in time, and how long it lasts, we need to follow three time-scales: hours, months, and years.

Hours and Days: The First Hit

Within twelve to twenty-four hours of the first oral dose of a broad-spectrum antibiotic, bacterial density in the colon begins to fall. By day three, in most adults, total faecal bacterial load has dropped by one to two orders of magnitude — that is, by a factor of ten to one hundred. Species diversity, measured by any of the standard ecological indices, drops sharply. The composition becomes *dominated* rather than *distributed*: a few surviving species, often organisms that were

present as minority members before treatment, expand into the empty niches and come to represent a large share of the total community. *Enterococcus* species are classic examples. So are certain *Enterobacteriaceae*, including the kinds of *E. coli* and *Klebsiella* that we would prefer not to see multiplying unchecked in the gut [REF:dethlefsen2011].

This early phase is also when patients notice the clinical side-effects. Antibiotic-associated diarrhoea affects somewhere between five and thirty per cent of patients, depending on the drug; for clindamycin and broad-spectrum beta-lactams, the rate is at the high end. Most cases are mild and self-limiting, caused by loss of the normal fermentative community that processes undigested carbohydrates in the colon. Without that community, carbohydrates reach the distal bowel intact, draw water osmotically, and produce diarrhoea. The same loss removes a protective function: the commensals that normally outcompete pathogens for nutrients and attachment sites are no longer there to do it.

The most severe version of this failure is *Clostridioides difficile* colitis — an infection that, almost uniquely among gut pathogens, is essentially a disease of antibiotic use. *C. difficile* spores are ingested from the environment all the time, but they cannot establish in a healthy, competitive gut community. When antibiotics remove the competition, the spores germinate, the bacterium colonises the damaged mucosa, and it releases toxins that cause severe inflammation and diarrhoea. In its worst form, *C. difficile* colitis produces life-threatening bowel perforation and requires emergency surgery. It kills approximately thirty thousand people per year in the United States alone, and almost every case is traceable to a prior course of antibiotics — often antibiotics prescribed for a trivial or questionable indication [REF:lessa2015]. The irony is stark: a drug prescribed to treat an infection sometimes causes one that is considerably worse.

Months: Partial Recovery

If the patient survives the acute phase — as the great majority do, with or without symptoms — what happens next is a long, gradual, partial return toward baseline. The Palleja study captured this phase in detail. Bacterial density recovers fairly quickly; within two to three weeks of the last dose, stool samples again look “full” under the microscope. Diversity recovers more slowly, usually taking one to two months. And composition — the precise identity and proportion of the species present — recovers most slowly of all, and often incompletely.

A different landmark paper, published by David Relman’s group at Stanford in 2011, followed three adults through two courses of ciprofloxacin, spaced six months apart. The first course reduced bacterial diversity substantially; by the time of the second course, the community had recovered — but not fully, and not to exactly the same configuration it had before. The second course produced a similar disruption, and the recovery from *that* course was less complete than the first. Each insult left a residue [REF:dethlefsen2011].

This raises an important point about how recovery should be measured. If we ask, “Did the bacterial community return to something that *looks* like a normal human microbiome?”, the answer is usually yes. The returning community has human-like numbers and human-like proportions of the dominant phyla. If we ask instead, “Did the community return to the specific configuration that existed in *this individual* before treatment?”, the answer is often no. The post-antibiotic community is a normal microbiome, but it is a *different* normal. It belongs to the individual’s post-treatment life, not to their pre-treatment life.

For a young, healthy adult with a single course of antibiotics for a real infection, this distinction may

be academic. The new microbiome is different, but probably functions well enough. For a child, or a chronically ill patient, or someone who has had many courses of antibiotics over a lifetime, the distinction matters a great deal.

Years: The Cumulative Shadow

The long-term consequences of antibiotic use — the ones that show up years or decades later — are the hardest to study and, not coincidentally, the ones Blaser has spent most of his career worrying about. The fundamental problem is that you cannot run a randomised controlled trial in which one group of children receives antibiotics and the other does not, and then follow both groups for thirty years to see who develops asthma, obesity, inflammatory bowel disease, or diabetes. The ethical and logistical barriers are insurmountable. What researchers can do, and have increasingly done, is look at very large observational cohorts — datasets containing millions of children, linked to national prescribing records and subsequent health outcomes — and ask whether the children who received antibiotics at particular ages went on to develop particular diseases more often than those who did not.

These studies cannot prove causation. They can only show association. An association between antibiotic exposure in the first year of life and a later diagnosis of asthma could, in principle, reflect a confounder: perhaps children who wheeze are more likely to be given antibiotics, and are also more likely to be diagnosed with asthma later. This kind of “reverse causation” is a real concern, and the best studies have gone to considerable lengths to address it — by excluding respiratory antibiotics from the analysis, by controlling for infection history, by comparing siblings who differ in their antibiotic exposure, and by applying causal inference techniques designed to strip out bias.

The associations that survive this scrutiny are uncomfortable. Large cohort studies have linked antibiotic exposure in the first two years of life to increased risks of childhood asthma, allergic rhinitis, atopic dermatitis, obesity, inflammatory bowel disease, juvenile idiopathic arthritis, and autism spectrum disorder — each association modest in magnitude, each of them repeatedly replicated in different populations [REF:aversa2021]. The effect sizes are not alarming for any one condition: a fifteen-to-thirty per cent increase in risk is typical, and a child who receives a single course of antibiotics for a genuine infection is not in a medical emergency. What is alarming is the cumulative picture, and the dose-response relationship. More courses produce more risk. Earlier exposure produces more risk than later. Broad-spectrum exposure produces more risk than narrow-spectrum. The pattern is exactly what you would expect if some fraction of the modern epidemic of childhood immune and metabolic disease were, in part, a consequence of ecosystem disruption during the developmental window when the microbiome is still being assembled.

This is the window we turn to next.

18.3 Early-Life Antibiotic Exposure and Lifetime Consequences

In Chapter 3, we described the assembly of the infant microbiome as a developmental process — a sequence of ecological successions that begins with the birth canal, continues through breastmilk, and settles into something resembling an adult community somewhere around the third birthday. This is one of the most consequential developmental processes in human biology. The microbes that establish themselves during the first thousand days educate the immune system, calibrate the

metabolism, and occupy the niches that will be hardest to reclaim later. Disrupt the process once, early enough and severely enough, and you may shape the rest of the life.

The Kenneth Ley Experiments

The strongest evidence for this claim comes from experiments that could only be done in animals. In 2012, a team led by Ilseung Cho and Martin Blaser at NYU reported a series of experiments on young mice given subtherapeutic doses of antibiotics — the kind of low-dose, continuous exposure used for decades in the livestock industry to promote growth. The mice were given penicillin, vancomycin, chlortetracycline, or a combination, in doses designed to mimic agricultural exposure rather than to treat any specific infection. The intervention began at weaning and continued for a few weeks [REF:cho2012].

The mice gained weight. More precisely, they developed altered body composition: more fat, different fat distribution, changes in hepatic lipid metabolism, altered expression of genes involved in lipid handling. The antibiotic-exposed animals were not sick — they looked like normal, healthy mice who happened to be slightly fatter than expected. The microbiome changes that accompanied these metabolic changes were modest in taxonomic terms but functionally significant: shifts in the proportions of the dominant phyla, changes in genes involved in carbohydrate metabolism and short-chain fatty acid production.

A follow-up study, published two years later, went further. The team showed that giving the antibiotic exposure *only* during the early-life window — the first few weeks after weaning — produced lasting metabolic changes, even though the microbiome eventually recovered its pre-antibiotic composition [REF:cox2014]. The gut community healed. The body did not. Something about the community's state during the early-life window had programmed the animal's later physiology, and once the window closed, the programming was set. You could restore the microbiome and still have the altered animal.

Transplant experiments confirmed the causal direction. When the researchers transplanted microbiota from antibiotic-exposed young mice into germ-free recipients, the recipients developed the same metabolic changes. When they transplanted microbiota from control mice, the recipients did not. The altered community was not just a marker of altered physiology; it was the cause of it. And the critical window — the period during which the alteration had to happen in order to produce the lasting effect — was early, brief, and easily missed.

The Human Parallel

Extrapolating from mice to humans is a dangerous business, and nowhere more dangerous than in microbiome research, where species differences in gut community structure, immune development, and metabolism are substantial. But the mouse experiments were designed to be informative about the human question, and their results lined up suspiciously well with what observational studies of children had been suggesting for years. Children given antibiotics before age two are, on average, modestly heavier than their peers later in childhood. Children given multiple courses of broad-spectrum antibiotics in the first year of life are at increased risk of obesity, metabolic disturbance, and immune disorders, as noted above. The human evidence is correlational, but the mouse evidence is experimental, and the two converge on the same conclusion: the developing microbiome is vulnerable in a way the adult microbiome is not, and the vulnerability has lifelong consequences.

One detail deserves emphasis. The agricultural practice that the mouse experiments modelled —

feeding low-dose antibiotics continuously to young livestock — was not adopted because someone had a theory about microbiomes. It was adopted empirically, in the 1950s, when farmers noticed that young chickens and pigs given antibiotic-laced feed grew faster and produced more meat per unit of feed. The practice spread worldwide and was, for decades, one of the main uses of antibiotics in volume terms: more antibiotic by weight went into livestock feed than into human medicine. We have, in other words, accidental empirical evidence from billions of animals that early-life, low-dose antibiotic exposure alters growth. Blaser's contribution was to point out the uncomfortable implication: if it does it in farm animals, there is no obvious reason it would not do it in children.

The Caesarean Complication

A full account of early-life microbiome disruption must also mention a procedure that is not usually thought of as an antibiotic exposure but effectively always includes one: the elective caesarean section. In most high-income countries, caesarean delivery is performed with prophylactic antibiotics given to the mother shortly before the skin incision, to reduce the risk of post-surgical wound infection. These antibiotics cross the placenta and reach the baby before the umbilical cord is cut. The first substance the newborn encounters is not the vaginal microbiome, but ampicillin or a cephalosporin.

We described the consequences of bypassing the birth-canal microbial transfer in Chapter 3: caesarean babies acquire a starter community dominated by skin and environmental organisms rather than the maternal vaginal and intestinal flora that vaginally-born babies inherit. Add to this the additional disruption of intrapartum antibiotic exposure, and the result is a newborn whose first microbiome is both incomplete and partially sterilised. The epidemiological signal is consistent with this mechanism: caesarean-born children show modestly elevated risks for many of the same conditions that are associated with early-life antibiotic use — asthma, allergy, obesity, immune disorders [REF:chu2017].

We do not yet know how to fully correct for this disruption. Experiments with “vaginal seeding” — swabbing the newborn with maternal vaginal fluid after caesarean delivery — have shown partial restoration of the expected microbial profile, but the clinical effects on later health outcomes are still being evaluated [REF:dominguezbello2016]. The point is not that caesarean sections should be avoided — they save lives, and in many cases are the only safe option for mother or baby. The point is that modern delivery practice has routinely imposed a double microbial insult on newborns without appreciating that either one was happening.

18.4 Antibiotic Resistance Genes in the Commensal Microbiome — A Hidden Reservoir

Most of what the public hears about antibiotic resistance concerns pathogens: MRSA in hospitals, resistant *E. coli* in urinary tract infections, multi-drug-resistant tuberculosis. The implicit mental model is that “resistance” is a property of “bad” bacteria — the ones that cause disease — and that the solution is to prescribe antibiotics more cautiously, develop new drugs, and contain the resistant pathogens in hospitals. This model is not wrong, exactly, but it is radically incomplete. It misses the fact that the largest reservoir of antibiotic resistance genes in most human bodies is not in the pathogens but in the commensals.

The Resistome

Microbiome researchers use the term *resistome* to describe the complete collection of antibiotic resistance genes in a community — all of them, whether in pathogenic or commensal bacteria, whether currently active or dormant, whether on chromosomes or on mobile genetic elements that can be transferred between cells. The resistome is, in some sense, a geological record of the community's antibiotic history: every gene in it is there because, at some point in the evolutionary past, an ancestor of the current community acquired it and found it useful enough to keep.

The surprise, when researchers first began characterising human gut resistomes systematically, was how large they were. Even in individuals with no recent antibiotic exposure, the gut microbiome carries hundreds to thousands of resistance genes, distributed across dozens of bacterial species [REF:hu2013]. Resistance to tetracyclines and beta-lactams is near-universal. Resistance to macrolides, aminoglycosides, and fluoroquinolones is common. Some of these genes are very old — older than the antibiotics we now use to treat human infections. The first antibiotic resistance genes evolved in the soil, where antibiotic-producing organisms have been making antibiotics, and competing organisms have been evolving resistance to them, for hundreds of millions of years. When humans began extracting antibiotics from soil organisms and using them as medicine, we imported, along with the drugs, the resistance infrastructure that had evolved alongside them.

What antibiotic treatment does to the resistome is not to *create* resistance genes — they are already there — but to *enrich* them. The sensitive bacteria die; the resistant bacteria survive and multiply. Genes that were present in one per cent of the community can be in thirty per cent a week later. And because many resistance genes sit on mobile genetic elements — plasmids, integrons, transposons, and the prophages we discussed in Chapter 16 — they can jump from one bacterial species to another. A resistance gene that was quietly present in a commensal *Bacteroides* can end up, after a few rounds of conjugation, in an opportunistic *E. coli* or a nascent *C. difficile* strain.

The Pool Contamination Problem

This dynamic produces an uncomfortable implication. Every course of antibiotics a patient takes not only treats the target infection but also updates the patient's personal resistome. The update is usually modest, but it is cumulative, and the gene pool of resistance genes is shared — through the family, through the workplace, through the healthcare system — with everyone the patient contacts. Antibiotic resistance, in this view, is not just a property of individual pathogens; it is a property of human populations.

A 2019 study by Yassour and colleagues followed Finnish children from birth to age three, sampling their gut microbiomes monthly and tracking their antibiotic exposures. The researchers found that resistance genes accumulated in the children's microbiomes over time, with each course of antibiotics producing a detectable uptick. Even more striking, some of the resistance genes detected in the children were genes for antibiotics they had never received — genes that had presumably been acquired by commensals through horizontal transfer, carried along as passengers, and then selected for when any antibiotic was given [REF:yassour2016]. The children were accumulating a resistance portfolio. They were not the patients for whom these resistance genes would matter, but they were the carriers who would pass them on.

This is the hidden cost of antibiotic use that traditional resistance metrics miss. When a hospital reports its rates of MRSA or VRE, it is counting clinically apparent resistant infections. It is not counting the slow accumulation of resistance determinants in the gut flora of every patient, every

visitor, every healthcare worker, and every patient's family at home. The hospital's resistance ledger looks manageable. The population's looks considerably less so.

18.5 *Candida* Blooms and the Kingdom That Fills the Vacuum

There is a group of organisms in the gut that we have discussed at length in this volume but have not yet mentioned in this chapter: the fungi. The oversight is appropriate, because for most healthy people, the gut mycobiome is a minor, well-regulated component of the microbial community — a minority by numbers, a minority by metabolic contribution, and normally kept in check by the bacterial majority. Antibiotics change that.

We touched on this phenomenon in Chapter 10, when we described how *Candida albicans* — the most common commensal yeast in the human gut — undergoes dramatic expansion after antibiotic treatment. The mechanism is straightforward: bacteria and fungi compete for the same resources — nutrients, attachment sites, and growth factors — and in a healthy gut, the bacteria usually win, because they are more numerous, they grow faster, and they produce metabolic by-products that suppress fungal proliferation. Remove the bacterial majority with a broad-spectrum antibiotic, and the fungi, which are intrinsically resistant to antibacterial drugs, find themselves in a newly empty environment with abundant nutrients and no competition. *Candida* populations in stool samples can increase by two or three orders of magnitude within a week of starting antibiotic treatment [REF:dollive2013].

For most patients, this fungal bloom is asymptomatic or produces only mild symptoms — oral thrush, vaginal candidiasis, diaper rash in infants. These are the visible consequences of an invisible shift in the gut ecosystem. For immunocompromised patients, the consequences can be much more serious, including invasive candidiasis and bloodstream infection. But the more interesting question, for the purposes of this chapter, is not what happens to the patient during the *Candida* bloom. It is what happens to the ecosystem, and what happens when the patient stops taking the antibiotic.

This is where the interkingdom dynamics we described in Chapter 10 come back into play. *Candida* is not a passive opportunist. Its expansion in the post-antibiotic gut actively remodels the environment in ways that make bacterial recolonisation harder. Fungal hyphae provide physical surfaces that can be colonised by opportunistic bacteria, creating mixed biofilms. *Candida* metabolism acidifies the local environment and produces signalling molecules that influence bacterial gene expression. Most importantly, *Candida* primes the host immune system in ways that favour inflammation — a pattern we discussed in Chapter 10 in the context of inflammatory bowel disease, where gut fungal expansion is a consistent feature. When the antibiotic is stopped and the bacterial community attempts to regrow, it must do so in an environment that has been reshaped by the intervening fungal bloom.

The clinical implications of this are still being worked out. It is clear, from observational studies, that patients with persistent gut *Candida* expansion after antibiotic treatment have worse microbiome recovery and higher rates of subsequent bacterial infections, including *C. difficile* [REF:markey2018]. It is also clear that the question of what to do about it is unresolved. Antifungal drugs such as fluconazole suppress *Candida* but introduce their own ecological disruption to the fungal community. *Saccharomyces boulardii*, a probiotic yeast we will return to in the probiotics chapter, has been used empirically to displace *Candida* and promote microbiome recovery, with

mixed but generally encouraging results. The idea that one fungus might be used as a defence against another's expansion is ecological medicine in its most literal sense.

For now, the key point is that the “kingdom blind spot” we discussed in Chapter 10 is not just a research problem. It is a clinical one. A clinician who prescribes antibiotics is implicitly prescribing a fungal expansion. The prescription note mentions neither.

18.6 The Recovery Phase — What Actually Happens When Antibiotics Stop

Imagine you could look into the gut of a patient who has just finished a course of antibiotics, with a time-lapse camera capable of following bacteria day by day. What you would see, over the following weeks, is one of the great dramas of ecological biology: a recolonisation race, played out in real time, with the outcome depending largely on who gets there first.

The Vulnerability Window

In the first week after the last dose, the gut community is at its minimum. Total bacterial density is low. Diversity is collapsed. The survivors — organisms that happened to be antibiotic-resistant, or that occupied sanctuary niches where the drug did not reach, or that existed as spores or other dormant forms — are scattered through an environment that is, from their perspective, unusually empty and unusually resource-rich. The normal competitive pressures that constrain any individual species in a healthy microbiome have been temporarily lifted.

This is the *vulnerability window*, and it lasts, depending on the antibiotic and the individual, somewhere between one and four weeks. During this time, the gut is abnormally permissive. Species that would normally struggle to establish, because they cannot compete with the incumbent community, can now find empty niches. Species that enter the gut in small numbers — from food, from the environment, from another person's skin — have a better chance of taking hold than they would in a healthy gut. The race is on.

The race has several competitors. There are the survivors of the antibiotic insult, regrowing from whatever remnant populations they managed to preserve. There are the organisms flowing in from upstream — from the mouth, from food, from the environment. There are the organisms in the healthcare environment, if the patient is in hospital, including some that are notorious for exploiting exactly this window: *Clostridioides difficile*, vancomycin-resistant *Enterococcus*, carbapenem-resistant *Enterobacteriaceae*. And, in a growing number of clinical studies, there are the organisms deliberately introduced by the clinician: probiotics, fermented foods, or — in the most radical version — transplanted faecal communities from a healthy donor.

What determines the outcome? Ecologically, the answer is: *who arrives first*, and in *what numbers*. This is the principle of *priority effects* in community ecology. When a habitat has been cleared and is being recolonised, early arrivers have a disproportionate influence on the final community structure, because they establish the conditions under which later arrivers must compete. A niche filled early by one species is not available to another. A metabolite produced early by one organism shapes the environment for everything that follows. The species that win the recolonisation race are not necessarily the ones that would dominate a mature, equilibrium community; they are the ones that happen to be abundant during the first days of the vulnerable window.

This insight has substantial therapeutic implications, which we will return to in the probiotics

chapter and in our discussion of faecal microbiota transplantation. For now, the point is that the recovery phase is not a passive process of the old community reassembling itself. It is an active, competitive, ecologically contingent process in which the starting conditions matter enormously. The “natural” recovery that happens when a patient stops their antibiotics and does nothing else is one possible outcome. It is not the only one.

Why Recovery Is So Often Incomplete

Return now to the Palleja study. Why did those nine bacterial species fail to recover at six months? Four possibilities, none of them reassuring, can be read from the data.

First, some species may have been *locally extinct*. If the antibiotic killed every cell of a given species in the gut, that species can only return if it re-enters from outside — from food, from contact with another person, from the environment. In an ancestral human living in close contact with dirt, family members, and livestock, re-entry would be relatively easy. In a modern adult living in a clean apartment and eating processed food, it may be nearly impossible. Extinction at the individual level becomes effectively permanent.

Second, some species may have depended on cooperative interactions with other species that were themselves eliminated. A bacterium that consumes the metabolic by-products of another bacterium cannot thrive if its partner is gone. A species that requires a specific vitamin produced by another species has the same problem. The gut microbiome contains extensive networks of such metabolic dependencies, and disrupting several species at once can cause cascading losses that are hard to reverse even after the antibiotic is withdrawn.

Third, some species may have been displaced by successor communities that now occupy the niches they used to hold. Once another organism has settled into a niche and established itself, the original occupant faces a competitive disadvantage in trying to return. The ecological door has closed.

And fourth, some species may still be present at very low, undetectable levels and may yet recover on a time-scale longer than the study measured. The six-month cut-off of the Palleja paper was dictated by practical constraints, not biological ones. We do not yet know the true time-scale of complete recovery — if it happens at all.

These four mechanisms are not mutually exclusive. They act together, and they act differently in different individuals, which is part of why recovery trajectories after antibiotic treatment are so variable. Some people seem to bounce back to essentially the same community they started with. Others do not. The factors that distinguish these groups — age, baseline diversity, diet, environmental exposure, genetics, and the specific antibiotic regimen — are still being untangled.

This is the background against which Chapter 19 opened with Sarah’s story. When Sarah received her third antibiotic course for recurrent urinary tract infection, her gut had not fully recovered from the second, and the community that regrew was progressively less able to provide the colonisation resistance that would normally keep opportunistic *E. coli* from proliferating. Each course had worked clinically, in the short term. Each course had left her ecosystem in a more vulnerable state than the one before. By the time the clinicians had run out of empiric options, Sarah’s problem was no longer an infection. It was an ecosystem that had been blunted by the very therapy meant to help it, and was now unable to defend itself.

18.7 Towards Precision: Narrow-Spectrum and Microbiome-Sparing Approaches

If the argument of this chapter is correct — that antibiotics, particularly broad-spectrum antibiotics, cause substantial collateral damage to the commensal microbiome, and that this damage has real consequences for patients — then it follows that the correct response is not to abandon antibiotics, but to use them better. The word “better” in this context means something more specific than “prescribing less.” It means matching the therapy more precisely to the target, so that the off-target damage is minimised.

This is not a new idea. The physicians who first used antibiotics in the 1940s understood the principle perfectly well: treat the specific bacterium causing the specific infection, with the most targeted drug available. What changed over the following decades was not the principle but the practice. As diagnostic laboratories became slower relative to clinical urgency, as culture-and-sensitivity results began to arrive too late to guide the initial prescription, and as broad-spectrum drugs became available that made the question of “which organism is it?” seem almost obsolete, empiric broad-spectrum therapy became the default. It was faster, easier, and “safer” in the sense that it was less likely to miss the pathogen. The ecological cost of that safety was invisible, so it was ignored.

Reversing this drift requires progress on several fronts.

Faster Diagnostics

The first front is diagnostic speed. The traditional workflow for bacterial identification — culture the specimen, wait one to three days for growth, test antibiotic susceptibility, wait another day for results — is too slow to guide the initial prescription in most acute infections. The physician has to start treatment before the laboratory can say what organism is present, and by the time the results are back, the empiric therapy is usually either working (in which case nobody bothers to change it) or not working (in which case switching is complicated by the partial treatment already received).

This workflow is finally changing. Rapid molecular diagnostics — PCR-based tests, mass spectrometry (MALDI-TOF), and metagenomic sequencing — can now identify bacterial pathogens from a clinical specimen in hours rather than days, and the same technologies can increasingly detect resistance genes directly, without waiting for the bacteria to grow. In some settings, a urine sample can be analysed for pathogen identity and antibiotic susceptibility in under an hour. In others, a bloodstream infection can be characterised overnight rather than over days [REF:goswami2023]. These advances have not yet reached most primary care settings, where the great majority of antibiotic prescriptions are written, but they are reaching there — slowly, patchily, and driven more by cost considerations than by microbiome concerns. The ecological consequences of faster diagnostics have been a pleasant side-effect rather than a primary motivation, but they are real. A prescription written against a known organism can use a narrower drug, administered for a shorter duration, than a prescription written against a statistical guess.

Narrow-Spectrum Drugs

The second front is the revival of interest in narrow-spectrum antibiotics — drugs that target specific bacterial families rather than the broad sweep that current broad-spectrum agents cover. For much of the antibiotic era, narrow-spectrum agents were considered commercially unattractive. A drug that works only against *Streptococcus pyogenes* has a smaller market than one that works

against streptococci, staphylococci, and a range of gram-negative bacilli. Pharmaceutical companies developing new antibiotics naturally preferred broader-spectrum candidates. The result was a pipeline biased toward exactly the kind of drugs that do the most collateral damage.

This bias is now being reconsidered, partly because broad-spectrum drugs are increasingly being held in reserve as last-line therapies in the face of resistance, and partly because microbiome research has made the collateral damage concrete. A narrow-spectrum antibiotic that leaves the bulk of the commensal community intact is now, for the first time, being valued for what it does *not* do. Drugs such as fidaxomicin — a narrow-spectrum macrocyclic antibiotic that targets *Clostridioides difficile* with minimal effect on the wider gut community — represent the kind of microbiome-sparing agent that was previously considered a curiosity and is now considered a template [REF:louie2011].

The logic of narrow-spectrum therapy is, however, only as good as the diagnostic precision that makes it possible. If you do not know what organism you are treating, a narrow-spectrum drug is a liability: it might miss. The diagnostic and pharmacological fronts advance together or not at all.

Shorter Courses

The third front is duration. For most infections, the duration of antibiotic therapy was set decades ago by expert opinion rather than by controlled trials. Ten days of penicillin for a streptococcal sore throat. Fourteen days of treatment for a urinary tract infection. Twenty-one days for a complicated pyelonephritis. These numbers were chosen to be safely long enough that the infection would be eradicated in everyone, including the slow responders. They were not chosen to minimise collateral damage, because at the time no one was thinking about that.

A growing body of modern trial evidence shows that shorter courses — sometimes *much* shorter courses — achieve the same clinical cures with less collateral damage. Three days of antibiotics for uncomplicated cystitis is now standard in many guidelines, replacing the traditional seven-to-ten-day course. Five days of antibiotics for community-acquired pneumonia produces outcomes indistinguishable from ten days. Seven days of treatment for bacteraemia is, in carefully selected patients, as effective as fourteen [REF:spellberg2016]. The pattern is consistent: the long courses that were standard for decades were longer than necessary, and the shortening delivers benefits both to the individual patient (fewer side-effects, less disruption) and to the community (less resistance selection). The guideline updates are happening slowly, because changing clinical practice is always slow, but the direction of travel is clear.

Protecting the Bystanders

The fourth front is the most speculative, but perhaps the most interesting. It asks: is there a way to protect the commensal microbiome *during* antibiotic therapy, so that the target infection is treated while the bystander community is shielded from collateral damage? The answers coming from research laboratories are beginning to include ideas that would have sounded absurd a decade ago.

One approach is to deliver antibiotics in a form that keeps them out of the colon. Oral beta-lactamases, taken alongside oral antibiotics, can degrade any drug residue that reaches the lower gut before it can do damage to the commensal community. A drug called ribaxamase was developed specifically for this purpose — to preserve the colonic microbiome during intravenous beta-lactam therapy — and early trials have shown that it substantially reduces the disruption associated with such treatment, including a reduction in *C. difficile* infection rates [REF:kokai2019]. The approach

is still experimental, but the principle — *selective decontamination of the contamination* — is novel and promising.

Another approach is to use bacteriophages, which brings us back to where this chapter began and forward to where the next chapter goes. Phages, as we have seen, are exquisitely specific. A phage that targets a particular strain of *E. coli* does not touch the rest of the community. In principle, phage therapy is the ultimate narrow-spectrum antibacterial treatment: it is a treatment for one strain of one species, and nothing else. The ecological consequences for the commensal community are, by design, minimal. We will explore what this means clinically in Chapter 19.

And a third approach, at the frontier of what is currently possible, is to combine antibiotic therapy with deliberate microbial reconstitution — either by giving probiotics during the antibiotic course to maintain a defensive population, or by giving faecal microbiota transplants afterwards to repopulate the damaged community. These are ideas we will develop in later chapters of this volume. For now, the point is that the mental model is shifting. For seventy years, antibiotic therapy was conceptualised as a duel between the drug and the pathogen, with the microbiome as inert background. That model is now being replaced by one in which the microbiome is an active, consequential participant — an ecosystem whose protection is part of the treatment plan rather than an afterthought.

A Different Kind of Stewardship

The standard response to the antibiotic crisis has been “antibiotic stewardship” — the set of institutional programmes designed to ensure that antibiotics are prescribed only when necessary, at the right dose, for the right duration, with the right drug. Stewardship is good. It is also, as currently practised, almost entirely concerned with preventing the emergence of resistant pathogens. The ecological damage to the commensal microbiome — the nine species that do not come back, the resistance genes that accumulate in healthy children’s guts, the metabolic shifts that persist for years — is not part of most stewardship programmes. It is not measured, not reported, and therefore not optimised against.

A broader form of stewardship is beginning to emerge. It treats the commensal microbiome as a patient-level asset worth protecting, measures antibiotic exposure in terms of ecological cost as well as pharmacological effect, and values narrow-spectrum, short-duration, microbiome-sparing therapy even when the broader-spectrum alternative would be clinically adequate. The call is not for less antibiotic prescribing but for *better* antibiotic prescribing: prescribing that accounts for the full cost of the treatment, including the cost to the organ that has, so far, been asked to bear that cost in silence.

Seventy-five years after Fleming’s Nobel Lecture, we are beginning to understand the dimensions of the warning he issued. It was not only that bacteria would evolve resistance to the drugs we use against them, although they have. It was also that those drugs, used the way they came to be used, would reshape the ecosystems we carry inside us — silently, cumulatively, and for keeps. The forest analogy is not an exaggeration. The cost, counted properly, is not negligible. And the question — what is a better way? — takes us directly into the next chapter, where we meet the oldest answer to that question, revived just in time.

Key References — Chapter 18

- [REF:fleming1945] Fleming A (1945). Penicillin. Nobel Lecture, 11 December 1945. Nobel Foundation, Stockholm.
- [REF:blaser2014] Blaser MJ (2014). *Missing Microbes: How the Overuse of Antibiotics Is Fueling Our Modern Plagues*. Henry Holt and Company, New York.
- [REF:palleja2018] Palleja A, Mikkelsen KH, Forslund SK et al. (2018). Recovery of gut microbiota of healthy adults following antibiotic exposure. *Nat. Microbiol.*, 3(11), 1255–1265. doi:10.1038/s41564-018-0257-9
- [REF:dethlefsen2011] Dethlefsen L & Relman DA (2011). Incomplete recovery and individualized responses of the human distal gut microbiota to repeated antibiotic perturbation. *Proc. Natl. Acad. Sci. USA*, 108(Suppl 1), 4554–4561. doi:10.1073/pnas.1000087107
- [REF:lessa2015] Lessa FC, Mu Y, Bamberg WM et al. (2015). Burden of *Clostridium difficile* infection in the United States. *N. Engl. J. Med.*, 372(9), 825–834. doi:10.1056/NEJMoa1408913
- [REF:aversa2021] Aversa Z, Atkinson EJ, Schafer MJ et al. (2021). Association of infant antibiotic exposure with childhood health outcomes. *Mayo Clin. Proc.*, 96(1), 66–77. doi:10.1016/j.mayocp.2020.07.019
- [REF:cho2012] Cho I, Yamanishi S, Cox L et al. (2012). Antibiotics in early life alter the murine colonic microbiome and adiposity. *Nature*, 488(7413), 621–626. doi:10.1038/nature11400
- [REF:cox2014] Cox LM, Yamanishi S, Sohn J et al. (2014). Altering the intestinal microbiota during a critical developmental window has lasting metabolic consequences. *Cell*, 158(4), 705–721. doi:10.1016/j.cell.2014.05.052
- [REF:chu2017] Chu DM, Ma J, Prince AL et al. (2017). Maturation of the infant microbiome community structure and function across multiple body sites and in relation to mode of delivery. *Nat. Med.*, 23(3), 314–326. doi:10.1038/nm.4272
- [REF:dominguezbello2016] Dominguez-Bello MG, De Jesus-Laboy KM, Shen N et al. (2016). Partial restoration of the microbiota of cesarean-born infants via vaginal microbial transfer. *Nat. Med.*, 22(3), 250–253. doi:10.1038/nm.4039
- [REF:hu2013] Hu Y, Yang X, Qin J et al. (2013). Metagenome-wide analysis of antibiotic resistance genes in a large cohort of human gut microbiota. *Nat. Commun.*, 4, 2151. doi:10.1038/ncomms3151
- [REF:yassour2016] Yassour M, Vatanen T, Siljander H et al. (2016). Natural history of the infant gut microbiome and impact of antibiotic treatment on bacterial strain diversity and stability. *Sci. Transl. Med.*, 8(343), 343ra81. doi:10.1126/scitranslmed.aad0917
- [REF:dollive2013] Dollive S, Chen YY, Grunberg S et al. (2013). Fungi of the murine gut: episodic variation and proliferation during antibiotic treatment. *PLoS ONE*, 8(8), e71806. doi:10.1371/journal.pone.0071806
- [REF:markey2018] Markey L, Shaban L, Green ER et al. (2018). Pre-colonization with the commensal fungus *Candida albicans* reduces murine susceptibility to *Clostridium difficile* infection. *Gut Microbes*, 9(6), 497–509. doi:10.1080/19490976.2018.1465158
- [REF:goswami2023] Goswami C, Fox S, Holden MTG & Connor TR (2023). Genomic approaches to rapid diagnosis of antimicrobial resistance. *Clin. Microbiol. Infect.*, 29(7), 841–847. doi:10.1016/j.cmi.2023.02.024
- [REF:louie2011] Louie TJ, Miller MA, Mullane KM et al. (2011). Fidaxomicin versus vancomycin for *Clostridium difficile* infection. *N. Engl. J. Med.*, 364(5), 422–431. doi:10.1056/NEJMoa0910812
- [REF:spellberg2016] Spellberg B (2016). The new antibiotic mantra — “shorter is better”.

JAMA Intern. Med., 176(9), 1254–1255. doi:10.1001/jamainternmed.2016.3646

- [REF:kokai2019] Kokai-Kun JF, Roberts T, Coughlin O et al. (2019). Use of ribaxamase (SYN-004), a beta-lactamase, to prevent *Clostridium difficile* infection in beta-lactam-treated patients: a double-blind, phase 2b, randomised placebo-controlled trial. *Lancet Infect. Dis.*, 19(5), 487–496. doi:10.1016/S1473-3099(18)30731-X

Chapter 19: Phage Therapy — The Comeback

19.1 A History: D'Hérelle, Soviet Medicine, and the Road Not Taken

In 1919, two years after he named the bacteriophage and published his discovery in the *Comptes Rendus* of the French Academy of Sciences, Félix d'Hérelle did something that no modern drug regulator would countenance. He drank a phage preparation himself, waited to see if anything happened, then gave it to his family members. When nobody fell ill, he administered phages to a twelve-year-old boy dying of severe dysentery at the Hôpital des Enfants-Malades in Paris. The boy recovered. D'Hérelle published the result and moved on to treating cholera patients in India, plague victims in Egypt, and dysentery outbreaks wherever he could find them. By the early 1920s, he was the most controversial microbiologist in the world — a self-taught outsider making extraordinary claims about a therapy that the scientific establishment could not decide was real [REF:dherelle1917].

We told part of this story in Chapter 16, when we described the discovery of phages and the bitter dispute between d'Hérelle and Jules Bordet over whether phages were truly living viruses or merely bacterial enzymes. But the story of phage *therapy* — the deliberate use of phages to treat human infections — has a trajectory of its own, and it is a trajectory shaped as much by geopolitics as by science. It is a story of a promising therapy that flourished in the East while the West forgot it existed, and of the extraordinary circumstances that are now bringing it back.

The Georgian Connection

The pivotal figure in the survival of phage therapy was a Georgian microbiologist named George Eliava. Born in 1892, Eliava trained in bacteriology in Tbilisi and Odessa before travelling to Paris in the early 1920s, where he met d'Hérelle at the Pasteur Institute. The two became collaborators and friends. Eliava was captivated by phage biology and by d'Hérelle's vision of therapeutic application. When he returned to Tbilisi, he carried that vision with him, and in 1923 he established a bacteriological research institute dedicated to studying and producing phages for clinical use [REF:kutateladze2010].

D'Hérelle visited Tbilisi twice, in 1933 and 1934, and spent approximately eighteen months working alongside Eliava — his longest sustained collaboration. Stalin himself had endorsed the project, recognising the strategic value of a therapy that could treat battlefield infections. The Tbilisi institute became the world's first purpose-built centre for phage research and production, and the two men began building a phage library: a collection of bacteriophages active against the most common human pathogens [REF:kutateladze2010].

The collaboration ended with a knock on the door. In 1937, during Stalin's Great Terror, George Eliava was arrested by the NKVD — Lavrenty Beria's secret police — and executed as an “enemy of the people.” The exact charge remains disputed; some accounts suggest political disfavour, others that Beria held a personal grudge. All of Eliava's papers and photographs were destroyed. He became a non-person, virtually unknown outside Georgia for forty years. D'Hérelle, devastated, returned to France and never went back to the Soviet Union [REF:kutateladze2010].

But the institute survived. Eliava's colleagues continued the work — quietly, under conditions of political repression, with whatever resources they could secure. Throughout the Soviet era, the Tbilisi institute (officially renamed the George Eliava Institute of Bacteriophage, Microbiology and Virology in 1988) produced phage preparations on an industrial scale. Soviet soldiers received phage cocktails as routine prophylaxis. Phage preparations were available in Georgian pharmacies without a prescription — as commonplace as aspirin. The science was real, the clinical use was widespread, and the results, while inconsistently documented by Western standards, were often impressive [REF:abedon2011].

Why the West Walked Away

Meanwhile, the West abandoned phage therapy almost entirely. The reasons were partly scientific and partly historical.

The scientific problems were genuine. Early phage preparations were often crude — poorly purified, inconsistently potent, and contaminated with bacterial debris that could cause fevers and inflammatory reactions. Clinical trials in the 1920s and 1930s produced mixed results, and the rigour of those trials, by modern standards, was poor. Some researchers used phages against infections that we now know were caused by organisms the phages could not infect. Others failed to account for the extreme specificity of phages — a preparation active against one strain of *Staphylococcus aureus* might be useless against the strain infecting a particular patient. Without proper strain matching, phage therapy was, in effect, a lottery.

But the decisive blow was not scientific. It was pharmacological. In 1928, Alexander Fleming discovered penicillin. By the early 1940s, industrial production of penicillin had begun, and the results were transformative. Penicillin worked against a broad range of bacteria. It did not require strain matching. It did not need to be individually prepared for each patient. It could be manufactured in vast quantities, shipped anywhere, stored indefinitely, and prescribed by any physician. It was, in every practical sense, the opposite of phage therapy: a universal solution where phages offered only a bespoke one.

The pharmaceutical industry chose antibiotics, and the funding followed. By the 1950s, phage therapy had virtually disappeared from Western medical literature. It persisted only in the Soviet Union and its satellite states — in Tbilisi, in Warsaw, in Wrocław — where it continued to be used clinically, studied scientifically, and refined practically, largely invisible to the English-speaking world.

The Cold War made this invisibility worse. Soviet-era publications were in Russian or Georgian, rarely translated, and published in journals that Western libraries did not carry. The few Western scientists who visited Tbilisi found an institute that was clinically productive but, after the Soviet collapse in 1991, physically crumbling. During the Georgian Civil War of the early 1990s, researchers reportedly stored phage cultures in their own homes — in personal refrigerators — to prevent the library from being lost. The institute survived, but just barely [REF:kutateladze2010].

The Return

The resurgence of phage therapy in the West was driven by a crisis that d'Hérelle could not have anticipated: the rise of antibiotic resistance. By the 2000s, the miracle drugs that had displaced phages were failing. Multi-drug-resistant bacteria — MRSA, carbapenem-resistant *Enterobacteriaceae*, extensively drug-resistant *Mycobacterium tuberculosis* — were causing infections that no available antibiotic could treat. The pipeline of new antibiotics had slowed to a trickle, because the

economics of antibiotic development are punishing: a drug that works is used for short courses and then held in reserve, producing far less revenue than a chronic-disease medication taken for years.

In this context, Western medicine began to look, with increasing urgency, at the therapy it had abandoned seventy years earlier. And when it looked, it found the Eliava Institute still standing — still producing phages, still treating patients, still maintaining a library of over six hundred characterised bacteriophages active against the most clinically important human pathogens [REF:zaldastanishvili2021].

19.2 Two Patients, Two Infections, Two Worlds

To understand what phage therapy actually looks like — not as a concept but as a clinical experience — it helps to follow two hypothetical patients through the same infection. Both have the same problem: a chronic, antibiotic-resistant urinary tract infection caused by multi-drug-resistant *Escherichia coli*, an infection that has recurred despite three courses of antibiotics over the past eighteen months. One patient follows the conventional Western pathway. The other travels to Tbilisi.

These are composite patients — not real individuals, but their stories are assembled from published case reports and documented clinical workflows at the Eliava Phage Therapy Center [REF:zaldastanishvili2021] [REF:leitner2021].

The Antibiotic Pathway: Sarah's Story

Sarah is fifty-two, a schoolteacher in Melbourne, and she has been battling recurrent UTIs for two years. The first infection responded to trimethoprim. The second did not — her GP switched to nitrofurantoin, which cleared the symptoms for three months before they returned. The third course was a fluoroquinolone, ciprofloxacin, which worked briefly but left Sarah with tendon pain and persistent gastrointestinal disturbance. Her urine culture now grows *E. coli* resistant to all three classes.

Sarah's experience is not unusual. Among patients with recurrent UTIs, resistance accumulates with each treatment cycle. A large cohort study of nearly 149,000 UTI patients found that by the sixth infection, sixty-five per cent of causative organisms were resistant to at least one antibiotic class, and twenty per cent were resistant to three or more [REF:dickstein2023]. Each course of empiric therapy — prescribed before culture results are available, targeting the most likely pathogen with the most likely drug — is a gamble. When the gamble fails, the patient has gained nothing except collateral damage to her microbiome.

And the collateral damage is real. A landmark 2018 study by Palleja and colleagues followed twelve healthy young men given a four-day course of three broad-spectrum antibiotics. Using deep metagenomic sequencing, the team tracked the gut microbiome's recovery over six months. The headline finding was that most bacterial species returned within about forty-five days. The less-publicised finding was that nine common species that were present before treatment remained undetectable in most subjects at the six-month mark. They had not recovered. They may never recover [REF:palleja2018].

Sarah's GP refers her to a urologist, who orders more cultures and prescribes a fourth antibiotic — meropenem, a carbapenem, one of the last-resort drugs. This requires intravenous infusion in a hospital outpatient unit. The infection clears. Three months later, it returns. The *E. coli* is now

carbapenem-intermediate. Sarah is running out of drugs, and each course has pushed her further down a narrowing corridor of resistance while progressively degrading the microbial community that, as we discussed in Chapter 5, normally provides colonisation resistance against the very pathogens now overwhelming her.

This is the paradox of empiric antibiotic therapy for chronic infections. The treatment is fast, accessible, and requires no specialised laboratory matching — any GP can write the prescription within minutes. But it is also blunt, ecologically destructive, and self-undermining: each course of broad-spectrum antibiotics selects for the resistant organisms that cause the next infection. For acute, first-episode infections in otherwise healthy people, the system works tolerably well. For chronic, recurrent, or resistant infections, it can feel like trying to put out a fire by pouring petrol on it.

The Phage Pathway: Elena's Story

Elena is fifty-four, a civil engineer from Munich, and she has the same infection — chronic recurrent UTI, multi-drug-resistant *E. coli*, three failed antibiotic courses. Her urologist, having exhausted the formulary, mentions a possibility she has never heard of: phage therapy at the Eliava Phage Therapy Center in Tbilisi, Georgia.

Elena's first step is not getting on a plane. It is getting on the phone. The Eliava center offers both in-person treatment and a distance protocol. She sends her most recent urine culture results to the center — crucially, including the live bacterial isolate, shipped in a transport medium via a medical courier service. This step is essential. Unlike antibiotics, which are prescribed based on the *class* of bacteria causing the infection, phage therapy is prescribed based on the *specific strain*. The phages that will treat Elena must be tested against *her* bacteria, in *her* culture, before a single dose is prepared [REF:zaldastanishvili2021].

At the Eliava laboratory, technicians plate Elena's *E. coli* isolate and test it against the center's standard phage preparations. The Eliava Institute maintains a library of commercial phage cocktails — products with names like Pyo Bacteriophage, Intesti Bacteriophage, and SES Bacteriophage — each containing a mixture of phages active against common pathogens. The testing method is conceptually simple: a drop of phage preparation is spotted onto a lawn of Elena's bacteria. If the phages kill the bacteria, a clear zone — a *plaque* — forms in the bacterial lawn. No plaque, no activity. Plaque means the phage can infect and lyse Elena's strain [REF:zaldastanishvili2021].

In Elena's case, one of the standard preparations shows good activity against her isolate. This is the best-case scenario. If none of the commercial products worked, the center would move to a second step: screening their broader research library of over six hundred characterised phages. If that also failed, they could prepare a *custom* phage — a phage specifically isolated and amplified to target Elena's strain, a process that takes eight to twelve weeks and costs significantly more [REF:zaldastanishvili2021].

Elena decides on in-person treatment and flies to Tbilisi. Her treatment course will last approximately ten to fourteen working days. On arrival, she is evaluated by a physician at the Phage Therapy Center — a clinical consultation that includes a full medical history, review of prior treatments, and fresh specimen collection to confirm the infecting organism has not changed since her samples were shipped.

The treatment itself is multi-route. Elena receives oral phage preparations — taken as a liquid, typically twice daily — for systemic exposure. Because her infection is in the urinary tract, she also receives phages delivered locally: vaginal suppositories containing the same phage preparation,

designed to establish high phage concentrations at the infection site. For prostatitis patients, the center uses rectal suppositories; for respiratory infections, nebulised inhalation. The route matches the anatomy of the infection, a principle that mirrors how any careful clinician chooses between oral and topical antibiotics — but taken further, because phages replicate at the site of infection, amplifying their own concentration precisely where the bacteria are [REF:zaldastanishvili2021] [REF:leitner2021].

What happens during the first week is, from Elena’s perspective, somewhat anticlimactic. There is no dramatic crisis, no intensive monitoring. She takes her oral doses, uses the suppositories, and attends periodic check-ups where fresh specimens are collected and cultured. The clinicians are watching for three things. First, is the bacterial load declining? Second, are any new organisms appearing — sometimes, as phages destroy one bacterial population, a previously suppressed species emerges from the biofilm underneath. Third, is the target strain developing phage resistance?

This third question is the critical one, and it is where phage therapy diverges most sharply from the antibiotic model. When bacteria develop antibiotic resistance, the usual response is to switch to a different antibiotic — often a broader-spectrum one, inflicting more collateral damage. When bacteria develop phage resistance, the response is to switch to a different phage — one that targets a different surface receptor on the bacterium. And here is the crucial insight from evolutionary biology: bacteria that evolve resistance to a specific phage often do so by modifying or losing the surface receptor that the phage uses for attachment. But surface receptors are not decorative. They serve essential functions — nutrient uptake, signalling, motility. Losing a receptor to escape a phage frequently comes at a fitness cost. In some well-documented cases, bacteria that evolve phage resistance simultaneously *lose* their antibiotic resistance, because the same surface structures are involved in both. The bacterium trades one vulnerability for another [REF:chan2016].

This phenomenon — sometimes called the “phage resistance trade-off” — is one of the most exciting features of phage therapy from an evolutionary perspective. It means that phage therapy and antibiotic therapy can, in principle, be used in combination to create an evolutionary trap: bacteria that resist the phage become sensitive to the antibiotic, and bacteria that resist the antibiotic remain susceptible to the phage. The organism cannot escape both pressures simultaneously.

Two Timelines, Two Philosophies

The contrast between Sarah’s and Elena’s experiences illustrates something deeper than a difference in treatment modality. It reflects two fundamentally different philosophies of infectious disease management.

The antibiotic model is *empiric*. It starts from probability: most UTIs are caused by *E. coli*, most *E. coli* are sensitive to trimethoprim, therefore prescribe trimethoprim. It is fast, cheap, and requires no microbiology beyond a standard culture (which, in many healthcare systems, is not even performed for a first uncomplicated UTI). It treats the infection as a generic category — a UTI — rather than as a specific organism in a specific patient. And it works well for first-episode acute infections, which is why it has been the standard of care for decades.

The phage model is *personalised*. It starts from the individual patient’s bacterium: isolate it, culture it, test the phages against it, select the ones that work, deliver them by the route that makes anatomical sense, monitor for resistance, adjust as needed. It is slower, more labour-intensive, and requires a level of microbiological engagement that most Western healthcare systems abandoned when antibiotics made it unnecessary. But for the patient whose infection has defeated the

empiric approach — the patient with a multi-drug-resistant chronic infection, the patient who has cycled through four or five antibiotic courses without lasting benefit — the personalised model offers something the empiric model cannot: a treatment matched to *their* infection, not to the statistical average.

The Georgian physicians at Eliava describe their approach as “iterative infectious disease management” — a phrase that captures the difference neatly [REF:zaldastanishvili2021]. Phage therapy is not a one-shot cure. It is a process: test, treat, culture, reassess, adjust. It is closer, in philosophy, to how oncologists manage cancer — with serial monitoring, adaptive treatment plans, and an acceptance that the disease may evolve under therapeutic pressure — than to how most infections are currently treated. It is slower and more demanding. But for the patients who need it, the question is not whether they can afford the time. It is whether they have any alternative.

19.3 Landmark Cases: When Compassionate Use Changed the Conversation

The case that brought phage therapy into the Western medical mainstream was not a clinical trial. It was a desperate gamble made by a wife who happened to be an epidemiologist.

The Patterson Case

In November 2015, Tom Patterson — a sixty-nine-year-old psychiatry professor at the University of California, San Diego — fell critically ill while on holiday in Egypt. He had developed a pancreatic pseudocyst, which became infected with *Acinetobacter baumannii*, a gram-negative bacterium notorious for its ability to acquire resistance genes. Patterson’s isolate was resistant to every available antibiotic. He was medevaced to Frankfurt, then transferred to UC San Diego, where he slipped into septic shock and multi-organ failure. His physicians had nothing left to offer [REF:schooley2017].

His wife, Steffanie Strathdee — a globally recognised HIV epidemiologist — began searching the literature. She found the phage therapy literature and, in a series of frantic emails and phone calls, contacted researchers across the United States. Within weeks, a collaboration had assembled that would have been impossible through normal channels. Teams at Texas A&M University, the United States Navy Medical Research Center, and San Diego State University screened their phage collections against Patterson’s *Acinetobacter* isolate. They found phages that worked. Forest Rohwer’s laboratory at SDSU purified the preparations for intravenous use [REF:schooley2017].

On March 15, 2016, phage therapy began — administered intravenously, directly into Patterson’s bloodstream. Four days later, he opened his eyes and recognised his daughter. It was the first sign of neurological recovery after weeks of unresponsiveness. By June, *Acinetobacter* was undetectable in his body. By August, he walked out of the hospital [REF:schooley2017].

The Patterson case was published in 2017 in *Antimicrobial Agents and Chemotherapy* and became the most widely reported phage therapy story in the Western media. It was not a controlled trial. It was not blinded. Patterson received other treatments alongside the phages. Scientifically, it proved very little in isolation. But it did something more important: it proved that phage therapy could be done — safely, rapidly, in a modern Western hospital, under FDA emergency investigational new drug (IND) authorisation — and that a patient at death’s door could walk out alive.

Strathdee and Patterson later co-authored a book, *The Perfect Predator*, and their case led directly to the founding of the Center for Innovative Phage Applications and Therapeutics (IPATH) at UC

San Diego — the first dedicated phage therapy centre in North America.

The Carnell-Holdaway Case

If the Patterson case demonstrated that phage therapy could work in an emergency, the case of Isabelle Carnell-Holdaway demonstrated something even more remarkable: that phages could be *engineered* for a patient.

In 2017, Isabelle — a fifteen-year-old British girl with cystic fibrosis — underwent bilateral lung transplantation at Great Ormond Street Hospital in London. Post-transplant, she developed a disseminated infection with *Mycobacterium abscessus*, an environmental mycobacterium that is intrinsically resistant to most antibiotics and notoriously difficult to treat. The infection spread to her surgical wound, her skin, and her liver. Her transplant team had exhausted all antibiotic options and, as Isabelle's condition deteriorated, her physicians contacted Graham Hatfull at the University of Pittsburgh [REF:dedrick2019].

Hatfull runs the largest collection of characterised mycobacteriophages in the world — the Science Education Alliance Phage Hunters Advancing Genomics and Evolutionary Science (SEA-PHAGES) programme, built over two decades with the help of thousands of undergraduate students. His team screened Isabelle's *M. abscessus* isolate against the collection and identified three candidate phages. One killed the bacteria efficiently. The other two did not — but Hatfull's team genetically engineered them, deleting the repressor gene that allowed the phages to enter lysogeny rather than killing their hosts, converting them from temperate to obligately lytic. The result was a bespoke three-phage cocktail designed specifically for Isabelle's infection [REF:dedrick2019].

Treatment began in June 2018, with intravenous infusions twice daily and topical application to skin lesions. Within seventy-two hours, the skin sores began to dry. After six weeks, the infection had substantially resolved. The case was published in *Nature Medicine* in 2019 and represented the first therapeutic use of genetically engineered phages in a human patient [REF:dedrick2019].

From Anecdotes to Evidence

These cases — and a growing number like them — are not clinical proof that phage therapy works. They are case reports, the lowest rung of the evidence ladder. The patients were desperately ill, received multiple concurrent treatments, and were not compared to matched controls. Any honest assessment must acknowledge that spontaneous improvement, placebo effect, or synergistic effects of co-administered antibiotics could account for some or all of the observed benefit.

But the case reports have accumulated to a point where the pattern is difficult to dismiss. A landmark 2024 retrospective analysed one hundred consecutive patients treated with personalised phage therapy across thirty-five hospitals in twelve countries. Clinical improvement was observed in 77.2 per cent. Bacterial eradication was achieved in 61.3 per cent. Adverse events were rare — seven non-serious suspected reactions in one hundred cases. Importantly, the study found that bacterial eradication was seventy per cent less likely when phages were given without concurrent antibiotics, suggesting that the combination of phages and antibiotics — the evolutionary trap we described earlier — may be more effective than either alone [REF:dedrick2023].

The evidence is not yet at the level that would satisfy a drug regulator for market approval. No phage therapy product has received full approval from the FDA, the European Medicines Agency, or the TGA. All current use in the West occurs under compassionate access pathways — the FDA's emergency IND mechanism, which can be approved by phone on the same day, or expanded access

protocols for individual patients. Approximately ninety phage therapy clinical trials are now active worldwide, including forty-one in the United States, and the first randomised controlled trials are beginning to report results [REF:fda2023].

But the gap between case reports and regulatory approval is being closed from both ends. From the bottom, the case series are growing larger and more systematically documented. From the top, regulators are developing new frameworks: Belgium has established a manufacturing monograph for phage preparations, Israel has created a “Clinical Phage Microbiology” pipeline analogous to clinical antibiotic susceptibility testing, and the FDA has issued guidance documents on phage therapy as a biological product [REF:fda2023] [REF:yerushalmy2023].

19.4 The Limits of the Comeback

It would be irresponsible to write about phage therapy without writing honestly about its limitations, because the history of medicine is littered with therapies that were oversold before they were understood.

The Georgian model at Eliava — the most mature clinical programme in the world — is itself instructive. The 2021 case series from the Eliava Phage Therapy Center described three patients with chronic bacterial infections treated with phages over extended periods. The results were nuanced, not triumphant. Patient 1, a forty-three-year-old man with cystic fibrosis and chronic *Pseudomonas aeruginosa* lung infection, was treated for four years. He improved — his bacterial load dropped by a factor of ten to one hundred — and he was able to discontinue antibiotics entirely, replacing them with phage monotherapy. But his *Pseudomonas* was never eradicated. Genetic analysis showed that the bacterial population had diversified under phage pressure, with sensitive and resistant strains coexisting. Patient 2, a sixty-four-year-old woman with bronchiectasis, underwent a similar trajectory: five sequential custom phages over three years, each replacing the last as resistance emerged. Her infection eventually cleared, but the path was long and iterative. Patient 3 — a seventy-two-year-old woman with chronic *Klebsiella pneumoniae* cystitis — did not respond. Her bacteria cleared initially but returned within months, and the treatment was ultimately unsuccessful [REF:zaldastanishvili2021].

These outcomes reflect the honest reality of phage therapy. It is not penicillin. It does not reliably eliminate an infection in a single course. It requires ongoing microbiological surveillance, willingness to switch phages as resistance emerges, and acceptance that some infections will not respond. For patients whose infections are otherwise untreatable, a seventy-seven per cent clinical improvement rate is substantial. But it is not a cure rate, and the twenty-three per cent who do not improve — and the thirty-nine per cent whose bacteria are not eradicated even when symptoms resolve — are a reminder that phage therapy is a sophisticated tool, not a magic bullet.

There are also practical barriers. Phage therapy requires microbiology infrastructure that many healthcare systems no longer maintain. It requires phage libraries — large, characterised collections — that take decades to build. It requires clinicians who understand phage biology, something that is not currently part of medical training anywhere in the Western world. And it requires a regulatory framework that does not yet exist for a therapy that is, by design, personalised: every patient potentially receives a different product, which does not fit neatly into the pharmaceutical model of standardised manufacturing and batch testing.

The comparison with Elena’s experience in Tbilisi is illuminating. Elena’s treatment cost approx-

imately €4,200 for a ten-to-fourteen-day in-person course. Whether that is cheap or expensive depends entirely on where you live. In the United States, where healthcare costs are among the highest in the world, a single hospital admission for intravenous antibiotics can easily exceed that amount many times over. In Australia, a public hospital might charge an uninsured foreign patient around AUD 1,300 per day — roughly comparable to the Eliava fee for a short IV course, but potentially exceeding it for a longer one. In many European countries with universal public healthcare, the patient's out-of-pocket cost for hospital-administered antibiotics might be close to zero. The point is not that phage therapy is cheap in absolute terms, but that the cost comparison depends heavily on geography and insurance — and that Elena is paying out of pocket for a treatment that no Western insurer currently covers. She also had to travel to Georgia, wait two to four weeks for an appointment, and entrust her care to a centre that, however experienced, operates outside the regulatory frameworks that Western patients and their physicians are accustomed to. For patients who cannot travel — the elderly, the immunocompromised, those in resource-limited settings — the Georgian model is not accessible. And for healthcare systems that measure quality by standardisation and reproducibility, a therapy that is different for every patient poses challenges that are as much administrative as scientific.

None of these limitations are reasons to dismiss phage therapy. They are reasons to invest in it — in clinical trials, in phage library development, in training programmes, in regulatory innovation. The antibiotic crisis is not going away. The pipeline of new antibiotic classes has produced only a handful of truly novel agents in the past four decades. The bacteria are evolving faster than the pharmaceutical industry can respond. Phage therapy is not the only alternative — antimicrobial peptides, lysins, CRISPR-based approaches, and microbiome engineering (which we will discuss in later chapters) all hold promise. But phage therapy has something the others do not: a century of clinical experience, a living proof-of-concept in Tbilisi, and a growing body of Western case data suggesting that, for the patients who need it most, the viruses that eat bacteria may be the oldest new medicine we have.

In the next section, we will turn to the cutting edge: the engineering of synthetic phages, the design of phage cocktails guided by genomics, and the question of whether phage therapy can ever scale from compassionate use to routine clinical practice.

Key References — Chapter 19

- [REF:dherelle1917] d'Hérelle F (1917). Sur un microbe invisible antagoniste des bacilles dysentériques. *Comptes Rendus Acad. Sci. Paris*, 165, 373–375.
- [REF:kutateladze2010] Kutateladze M & Adamia R (2010). Bacteriophages as potential new therapeutics to replace or supplement antibiotics. *Trends Biotechnol.*, 28(12), 591–595. doi:10.1016/j.tibtech.2010.08.001
- [REF:abedon2011] Abedon ST, Kuhl SJ, Blasdel BG & Kutter EM (2011). Phage treatment of human infections. *Bacteriophage*, 1(2), 66–85. doi:10.4161/bact.1.2.15845
- [REF:chan2016] Chan BK, Siström M, Wertz JE, Kortright KE, Narayan D & Turner PE (2016). Phage selection restores antibiotic sensitivity in MDR *Pseudomonas aeruginosa*. *Sci. Rep.*, 6, 26717. doi:10.1038/srep26717
- [REF:schooley2017] Schooley RT, Biswas B, Gill JJ et al. (2017). Development and use of personalized bacteriophage-based therapeutic cocktails to treat a patient with a dissemin-

- ated resistant *Acinetobacter baumannii* infection. *Antimicrob. Agents Chemother.*, 61(10), e00954-17. doi:10.1128/AAC.00954-17
- [REF:palleja2018] Palleja A, Mikkelsen KH, Forslund SK et al. (2018). Recovery of gut microbiota of healthy adults following antibiotic exposure. *Nat. Microbiol.*, 3(11), 1255–1265. doi:10.1038/s41564-018-0257-9
 - [REF:dedrick2019] Dedrick RM, Guerrero-Bustamante CA, Garlena RA et al. (2019). Engineered bacteriophages for treatment of a patient with a disseminated drug-resistant *Mycobacterium abscessus*. *Nat. Med.*, 25(5), 730–733. doi:10.1038/s41591-019-0437-z
 - [REF:zaldastanishvili2021] Zaldastanishvili E, Leshkasheli L, Dadiani M et al. (2021). Phage Therapy Experience at the Eliava Phage Therapy Center: Three Cases of Bacterial Persistence. *Viruses*, 13(10), 1901. doi:10.3390/v13101901
 - [REF:dickstein2023] Dickstein Y, Jeong JY, Brintz BJ et al. (2023). Antibiotic resistance in urinary tract infections: results from a multi-site cohort. *J. Infect. Dis.*, 230(6), e1344. doi:10.1093/jid/jiad123
 - [REF:dedrick2023] Dedrick RM, Smith BE, Cristinziano M et al. (2023). Personalized bacteriophage therapy outcomes for 100 consecutive cases: a multicentre, multinational, retrospective, observational study. *Nat. Microbiol.* doi:10.1038/s41564-024-01705-x
 - [REF:leitner2021] Leitner L, Ujmajuridze A, Chanishvili N et al. (2021). Intravesical bacteriophages for treating urinary tract infections in patients undergoing transurethral resection of the prostate: a randomised, placebo-controlled, double-blind clinical trial. *Lancet Infect. Dis.*, 21(3), 427–436. doi:10.1016/S1473-3099(20)30330-3 [Published online 2020; print March 2021]
 - [REF:fda2023] U.S. Food and Drug Administration (2023). Science and Regulation of Bacteriophage Therapy. FDA Guidance Document.
 - [REF:yerushalmy2023] Yerushalmy O, Khalifa L, Copenhagen-Glazer S et al. (2023). Towards Standardization of Phage Susceptibility Testing: The Israeli Phage Therapy Center “Clinical Phage Microbiology” — A Pipeline Proposal. *Clin. Infect. Dis.*, 77(Suppl 5), S337–S351. doi:10.1093/cid/ciad514
 - [REF:aswani2021] Aswani VH & Shukla SK (2021). An Early History of Phage Therapy in the United States: Is It Time to Reconsider? *Clin. Med. Res.*, 19(2), 82–89. doi:10.3121/cmr.2021.1605

Chapter 20: Diet — Feeding the Ecosystem

20.1 The Dry Season in Hadzaland

In the acacia woodlands south of Lake Eyasi, in northern Tanzania, a group of people still live in a way that most of humanity abandoned about ten thousand years ago. The Hadza are among the last full-time hunter-gatherers on earth. There are perhaps a thousand of them, and only a few hundred still live almost entirely on what they can hunt, dig, and pick from the landscape around them. They do not keep animals. They do not plant crops. They do not have a word for “farmer.” What they eat today is what they found today.

In 2013 and 2014, a team of researchers from Stanford University and the University of Bologna went to live with the Hadza and do something that had never been done before: follow the same people, through an entire year, and track what happened inside them as the seasons turned. They collected stool samples in the wet season, when the women’s foraging brought back baobab pods, berries, and tubers in abundance, and the men returned with honey stolen from wild bee hives high in the acacia trees. Then they collected samples again in the dry season, when the berries were gone, the tubers were woody and bitter, and the camp subsisted mostly on meat and *ekwa*, a hard root dug out of the cracked earth with pointed sticks [REF:smits2017].

What they found was astonishing. The gut microbes of the Hadza did not stay the same. They *cycled*. Entire families of bacteria that were abundant in the wet season almost disappeared in the dry season, and then came back when the rains returned and the fruits and tubers followed. One group, the Bacteroidetes — a family we have met many times in this book — surged when the diet was high in plant fibre and retreated when it was not. Another group, the Firmicutes, moved in the opposite direction. The researchers were watching, in real time, a rainforest of bacteria change species composition with the weather.

Now compare that with the microbes of a person in Melbourne, Manchester, or Minneapolis. Here the seasons, such as they are, do not reach the supermarket. A tomato in January tastes the same as a tomato in July, only worse. The diet does not cycle. The microbes do not cycle. Or rather — and this is the more disturbing finding from a different branch of the same research program — they have already stopped cycling. Many of the taxa that cycle in the Hadza are entirely absent from the Western gut. We are not in a permanent summer. We are in a permanent winter so long that most of us no longer remember anything else [REF:sonnenburg2016] [REF:schnorr2014].

This chapter is about that forgotten summer — and about how, three times a day, you get a vote on whether it returns.

20.2 The Gardener’s Dilemma

In Chapter 18 we talked about antibiotics as something closer to a forest fire than to a scalpel — a tool that treats the infection and sterilises the ground around it. In Chapter 19 we talked about

phages, which act more like a sniper than a flame-thrower. Both chapters were about killing things. This chapter is about feeding them.

The distinction matters because almost everything the popular science shelf of a health food store sells you is framed as a killing or a planting intervention. “Cleanses” kill. “Probiotics” plant. Both framings are, for most people most of the time, the wrong way to think about the problem. Your microbiome is not a garden in the sense that you start from bare earth and decide what goes where. By the time you are old enough to read a label on a supplement bottle, the garden is decades old. Some of your microbial species have been with you since you came down the birth canal. Others joined when you were weaned, or when you went to primary school, or when your family moved house. Many of them cannot be bought in a capsule at all. Some of them are probably extinct in most of your neighbours, and you will not get them back by swallowing a pill.

What you *can* do, every day, is decide what to feed them.

The metaphor to hold in mind for the rest of this chapter is not pharmacology. It is agriculture. A gardener does not usually plant a new lawn every spring; she feeds the one she has. She chooses composts and mulches that favour certain plants and disadvantage others. She waters in patterns that support the soil she wants, and she knows that what she cannot see — the fungi and bacteria and protists underneath — is doing more of the work than the bits she can. Over years, her choices produce a particular kind of garden. Somebody else, making different choices with the same starting soil, produces a different one. Neither of them got to pick the seeds.

This chapter is a practical guide to feeding the garden you have. It is organised around a question the reader is entitled to ask at the end of any chapter on diet: *so what should I eat on Tuesday?* We will get to a direct answer, tiered by the strength of the evidence, in section 20.12. Before we get there, we have to understand what we are feeding, what the microbes do with what we give them, and why the answers are not always the same for every gardener.

We will work outwards from the strongest and cleanest evidence — fibre — to progressively more complex and contested territory: fermented foods, polyphenols, the Mediterranean pattern, protein and fat, ultra-processed food, meal timing, personalisation, and the controversies (gluten, FODMAPs, lectins, carnivore) that readers will have encountered on social media and are entitled to have addressed directly. Nothing in this chapter is invented. Every claim is tied to a published study, and where the evidence is weak or preliminary I will say so plainly.

One warning before we begin. The history of nutrition science is littered with confident claims that did not survive the next decade. Eggs were bad, then good, then bad, then neither. Fat was the enemy, then sugar was the enemy, then insulin, then inflammation. A reasonable reader, watching that parade go by, might conclude that nutrition science is simply unreliable and the sensible thing to do is ignore it. That conclusion would be wrong, but it would be understandable. What I can promise is this: the evidence we will draw on in this chapter is *microbiome-mediated* evidence, which is a different kind of evidence from the older generation of “association studies” that drove the older reversals. We can now watch, in human subjects, what happens to specific microbial populations, specific metabolites, specific inflammatory markers, and in some cases specific clinical outcomes when the diet changes. The evidence is not perfect. But it is less likely to reverse, because it is grounded in a mechanism that can be observed directly.

With that, let us look at the most important thing any human being can put on the plate for the ecosystem inside them: fibre.

20.3 The Currency of the Realm: Fibre and the Microbiota-Accessible Carbohydrates

What your body cannot digest, the microbes can

When you eat a piece of wholemeal toast, your digestive enzymes break down most of it. The starches in the flour are split into glucose and absorbed in the small intestine. The proteins are cleaved into amino acids. The fats, if there are any, are emulsified by bile and picked up by the enterocytes that line the small bowel. By the time what remains of the toast reaches the large intestine, the components that your own enzymes know how to handle are gone.

What is left is the part of the plant that your body cannot digest. The botanical name for most of it is dietary fibre. The chemical reality is more interesting. “Fibre” is not one molecule. It is a bag of very different polysaccharides — cellulose from plant cell walls, hemicelluloses from the matrix that holds those walls together, pectins from the soft tissues of fruits, inulin from the roots of chicory and the bulbs of onions, resistant starch from cooked-and-cooled potatoes and unripe bananas, beta-glucans from oats and barley, arabinoxylans from wheat bran. Some of these you have heard of. Most of them you haven’t. What they have in common is that the human genome does not encode the enzymes to take them apart.

The microbes in your large intestine do. This is one of the most important biochemical facts in this entire book. Your genome carries instructions for making roughly seventeen digestive enzymes that act on carbohydrates. The combined genome of the microbes in your colon encodes tens of thousands of such enzymes, organised into complex multi-protein machines called carbohydrate-active enzymes, or CAZymes. The gut is, in effect, a second digestive system — one that runs on different software than your own, and that is capable of metabolising substrates your own system cannot touch [REF:koh2016].

In 2014, Erica and Justin Sonnenburg at Stanford proposed a name for the subset of dietary carbohydrates that reach the large intestine and feed this second digestive system: microbiota-accessible carbohydrates, or MACs [REF:sonnenburg2014]. The name sounds technical, but the concept is the single most important idea in modern dietary microbiology, and once you see it you cannot unsee it. Every meal you eat delivers some amount of MAC to your colon. The higher the MAC content of your meal, the more substrate your microbes have to work with. The lower the MAC content, the more your microbes have to find something else to eat — and as we are about to see, they do not starve quietly.

The mucus is also food

The human large intestine is lined with a layer of mucus that is not decorative. It is a physical and chemical barrier, two layers thick in the colon, that keeps the trillions of resident bacteria in the lumen from touching the epithelial cells underneath. The outer layer is loose and colonised by bacteria that graze on it. The inner layer is dense, highly cross-linked, and — in health — virtually sterile. The cells in the epithelium sit under a shield of what is, for the bacteria that grow above it, the last line between them and host tissue. This matters enormously because the moment bacteria reach the epithelium, the immune system notices, and inflammation follows.

In 2016, a team led by Mahesh Desai, working in the laboratory of Eric Martens at Michigan, put gnotobiotic mice — mice with a defined, human-derived gut community — through a controlled dietary experiment. One group got a diet containing plenty of plant polysaccharides. The other got a diet essentially devoid of fibre but nutritionally complete in every other respect. The team then

used fluorescent probes to measure the thickness and composition of the colonic mucus layer, along with metagenomics to see what the bacteria were doing [REF:desai2016].

The finding was so vivid it made the cover of *Cell*. On the fibre-free diet, the mucus layer thinned dramatically. The bacteria, unable to find the plant polysaccharides they were equipped to degrade, switched their metabolism to the only remaining source of carbohydrate in the bowel — the glycans of the host mucus layer itself. Specific species that carry the enzymatic machinery for mucin degradation, including *Akkermansia muciniphila* and *Bacteroides caccae*, bloomed. The bacteria were literally eating the wallpaper. Worse, when the mice were then challenged with a pathogen — *Citrobacter rodentium*, a rough mouse equivalent of enterohaemorrhagic *E. coli* — the fibre-deprived animals developed more severe disease than the fibre-fed ones. The thinned mucus layer had become a highway rather than a barrier.

Think about what that finding means. If you stop feeding your microbes fibre, they don't quietly shrink away. They pivot — and start eating *you*. The mechanism is not unique to mice. Mucin-degrading bacteria are abundant in the human colon, and the machinery they use to take mucus apart is, in biochemical terms, general-purpose. The experiment is hard to replicate in humans for ethical reasons, but the implication is serious: the thickness of the mucus layer that separates our bacteria from our own cells is, in part, a function of how much fibre we eat.

What your grandchildren will inherit

A second Sonnenburg paper, also in 2016, made the stakes even clearer. Using the same kind of mice carrying a human-derived microbiota, the team put animals on a low-MAC diet for a single generation and measured what happened to the next [REF:sonnenburg2016].

The parents lost some of their bacterial diversity on the low-MAC diet — roughly a 60% reduction in the species richness of their gut community. When they were put back on the high-MAC diet, most of what they had lost came back. Not all of it, but most. But when the mice were bred, the pups inherited only the species that had survived the low-MAC period. The pups never encountered the missing species and therefore never acquired them. By the fourth generation on the cycling diet, many of those species were gone entirely. Even when the grandchildren and great-grandchildren were put back on a high-MAC diet, the species did not come back. They were extinct.

The experiment is in mice, and we should be careful about what it predicts for humans. But the logic of extinction applies to any transmitted microbial community. If a mother carries species A, B, and C, and her child is born and breastfed, the child acquires species A, B, and C. If the mother has lost species C entirely before conception, the child cannot inherit what is not there. Multiple generations of low-fibre eating in a population can, in principle, produce exactly the pattern Sonnenburg's mice showed: a cumulative, generation-over-generation loss of diversity that does not bounce back even when diet improves, because the species simply are no longer in the transmission pool.

Look at the Western microbiome. Compared with the microbiomes of traditional populations — the Hadza, the Yanomami of the Amazon, the Matsigenka of the Peruvian jungle, rural Malawian children — it is strikingly impoverished. Between a third and a half of the bacterial species that those populations carry are simply not present in people who grew up in industrialised cities. We used to think the explanation was environmental exposure: that hunter-gatherers pick up more microbes from dirt, water, and animals than city dwellers do. The Sonnenburg mouse experiment adds another explanation, one that is harder to hear. We may have eaten some of them into extinction, across generations, without realising it [REF:sonnenburg2016].

This is why fibre deserves its place at the start of this chapter. It is not, as some decades of nutritional advice framed it, merely useful for keeping things moving. It is the substrate that keeps the entire bacterial community alive and operational in its normal configuration. Stop supplying it, and the community does not pause politely. It eats the house, loses members, and in the long run — if the pattern is maintained across generations — it loses them permanently.

What the microbes give back: short-chain fatty acids

When microbes ferment MACs — and “fermentation” here means anaerobic breakdown by bacterial enzymes, not the pickle-jar kind, which we will get to in section 20.6 — they produce a small number of remarkable molecules called short-chain fatty acids, or SCFAs. Three of them dominate: acetate (two carbons), propionate (three carbons), and butyrate (four carbons). They are produced in roughly a 60:20:20 ratio in a healthy adult colon, and between them they account for somewhere between 5 and 10 percent of the total daily energy intake of the host [REF:koh2016].

Each has a distinct job.

Butyrate is the most remarkable, because it is the primary energy source for the cells lining your own large intestine. The colonocytes — the epithelial cells of the colon — get roughly 70% of their energy from bacterial butyrate rather than from the blood supply. This is one of the most intimate metabolic relationships between any two kingdoms of life. You are not running your own colon on your own fuel. The bacteria that sit in the lumen are, quite literally, feeding the cells that face them. When butyrate production drops — because MAC intake drops, because the fibre-fermenting bacteria are not there, because antibiotics have killed them — the colonocytes fall back on less favourable fuels, and the tight junctions between them begin to loosen. The permeability of the gut wall, in other words, is partly a function of how well-fed the colonocytes are, and the colonocytes’ feeding is partly a function of how much fibre you ate yesterday.

Butyrate does more than feed cells. Inside the colonocyte, it acts as a signalling molecule. It inhibits an enzyme called histone deacetylase, or HDAC, which changes how certain genes are read. One of the downstream effects is the differentiation of a particular class of immune cells — regulatory T cells, or Tregs — that keep the colonic immune system from over-reacting to the bacteria it lives with. In 2013, two separate groups published almost simultaneously in *Nature* and in *Science* showing that butyrate drives Treg differentiation in the colon and suppresses colitis in mouse models [REF:furusawa2013] [REF:smith2013]. The food you eat, via the microbes that eat it, signals to the cells that decide whether your immune system tolerates what it sees in your gut or attacks it.

Propionate has its own story. Much of what the bacteria produce is absorbed into the portal vein and taken directly to the liver, where it influences gluconeogenesis — the making of new glucose — and lipid synthesis. Propionate appears to reduce appetite in some studies by acting on gut hormones like PYY and GLP-1, the same hormones targeted by the latest generation of weight-loss drugs. It is not a drug. But it is an endogenous molecule produced, by your bacteria, in proportion to the fibre you feed them.

Acetate is the most abundant and the most diffusible. It reaches the systemic circulation and is taken up by peripheral tissues, including muscle and adipose tissue. Its signalling role is less well understood than butyrate’s, but it contributes to the cross-talk between the gut and brain and appears to play a role in appetite regulation.

All three of them — acetate, propionate, butyrate — act, collectively, as chemical messengers that your microbes send to your cells. They are the receipt your microbes give you in exchange for the

fibre you supply. Low fibre in, low SCFAs out. And low SCFAs out means a colon that is both hungrier and more inflamed than it should be.

The hard outcomes

If the mechanism were all we had, a cautious reader would be right to ask whether it matters clinically. Cells in petri dishes do all sorts of things they do not do in living people, and mouse experiments do not always translate to the human bowel. The honest answer is that we also have clinical outcome data, and it is striking.

In 2019, a consortium led by Andrew Reynolds at the University of Otago in New Zealand published, in *The Lancet*, a series of systematic reviews and meta-analyses pooling data from 185 prospective cohort studies and 58 randomised controlled trials — nearly 135 million person-years of observation [REF:reynolds2019]. The question they asked was simple: does carbohydrate quality, and specifically dietary fibre intake, predict hard clinical outcomes?

The answer was yes, and the magnitude was larger than most people realise. For every eight grams per day of additional fibre intake, the analysis found an 8% reduction in all-cause mortality, a 7% reduction in coronary heart disease, a 19% reduction in colorectal cancer, and a 15% reduction in type 2 diabetes. The dose-response was continuous: more fibre, less disease, all the way up to the highest intake categories — about 25 to 29 grams per day — with no plateau in the data at conventional intake levels and suggestions that higher might be better still.

For context, the average daily fibre intake in most Western countries is around 12 to 15 grams. The recommendations from most national dietary guidelines are 25 to 30 grams per day for adults. The intake of traditional populations has been estimated at around 100 grams per day — roughly seven times what most of us manage. The Hadza, in the wet season, eat around 150 to 200 grams of fibre per day, because baobab pods and wild tubers are essentially concentrated fibre delivery devices. You would not enjoy eating what the Hadza eat. But the gap between their intake and ours is the gap between a microbiome that cycles through seasons and one that has forgotten what a season is.

The Reynolds analysis is not the only evidence, but it is the cleanest recent summary, and it gives us a number to carry forward: eight grams of fibre per day buys you roughly an 8% reduction in your risk of dying of anything. That is an extraordinary number for any dietary variable, and it is backed by the largest and most rigorous synthesis in the field.

The many shapes of fibre

One more thing before we move on, because “just eat more fibre” is not a useful instruction without some feel for what the word actually means.

Fibre is not one thing. The commercially useful distinction is between *soluble* and *insoluble* fibre, based on whether the fibre dissolves in water, but this distinction is a legacy of food chemistry, not microbial biology, and it maps only loosely onto what the microbes actually care about. A more useful breakdown, for our purposes, is functional:

Fermentable fibres are the ones the microbes can take apart readily and use as fuel. Inulin (in chicory root, Jerusalem artichoke, garlic, onion, leek, asparagus), resistant starch (in cooked-and-cooled potatoes, cooked-and-cooled rice, green bananas, and legumes), beta-glucans (in oats and barley), and pectins (in apples, citrus peel, and most fruits) fall into this category. These are the ones that generate the most SCFAs and feed the most microbes.

Slowly fermentable fibres — arabinoxylans in wheat bran, certain hemicelluloses in whole grains — are broken down more gradually and tend to be fermented further along the colon, producing SCFAs in regions where the first category has mostly been consumed. This matters because cancer risk in the colon is highest in the distal portion, and SCFA production in that area depends on slowly fermentable fibres that survive the early part of the journey.

Poorly fermentable fibres — cellulose and lignin, the structural fibres of plant cell walls and wood — are largely excreted intact. They have important non-microbial effects on bowel transit and stool bulk. But they do not, as such, feed the ecosystem.

A good diet contains all three classes. If you eat a bowl of oats topped with berries and a spoon of ground flax in the morning; some lentils or chickpeas and leafy greens at lunch; and a dinner of whole-grain pasta with vegetables, herbs, olive oil, and a small piece of fish, you have covered the full spectrum without thinking about it. If you eat white toast, a ham sandwich on white bread, and pasta with tomato sauce, you have not.

The 2024 review by Deehan and colleagues is the best single summary of fibre subtypes and their metabolic effects currently available, and readers who want the biochemical detail can find it there [REF:deehan2024]. For our purposes, the single most important take-away from the fibre section of this chapter is this: your microbes run on substrate, the substrate is fibre, and there are a great many different fibres with subtly different effects. The next section turns that insight into an actionable rule — a single number that, of everything in this chapter, is probably the most concretely useful thing the modern microbiome literature has produced.

20.4 Diversity Beats Dose: The Thirty-Plant Rule

A number you can actually use

If this chapter had to offer the reader one number, and only one, the number would be thirty.

It comes from the American Gut Project, a citizen-science study that, between 2012 and 2017, enrolled more than 11,000 volunteers in the United States, the United Kingdom, and Australia. Participants paid a small fee, were mailed a swab kit, sampled their own stool, and filled in a long dietary questionnaire. The samples were sequenced at the laboratory of Rob Knight at the University of California San Diego. The scale of the project — roughly an order of magnitude larger than any comparable microbiome study at the time — allowed the analysts to ask questions that smaller studies cannot answer, including: which dietary variable, out of dozens tracked, best predicts the diversity of the gut microbiome? [REF:mcdonald2018]

The answer was not total fibre grams. It was not whether the participant identified as vegan or vegetarian. It was not meat consumption, carbohydrate consumption, or calorie intake. It was a single variable the research team called “plant types per week”: the number of distinct plant species a person reported eating over the previous seven days.

The result fell cleanly into two groups. People who reported eating thirty or more different types of plants per week had, on average, significantly more diverse gut microbiomes than people who reported ten or fewer. The effect was dose-dependent, consistent across the three countries, and robust to adjustment for fibre intake, age, body mass index, and whether the person called themselves a plant-eater. Eating more *species* of plants, regardless of how much fibre was in any single serving, predicted a richer ecosystem in the gut.

Thirty plant species a week sounds like a lot. Most people, asked to guess how many they eat, guess something like ten or fifteen. Almost everybody underestimates. The trick is what you count. Every plant counts: fruits, vegetables, whole grains, legumes, nuts, seeds, herbs, and spices. A salad with mixed greens, tomato, cucumber, red onion, parsley, and a dressing of olive oil, lemon, and black pepper contains eight different plants before you have finished chewing. A bowl of muesli with oats, wheat flakes, barley, almonds, walnuts, sunflower seeds, flaxseed, dried cranberries, and raisins is nine. A good Indian dal with lentils, tomato, onion, garlic, ginger, turmeric, cumin, coriander, chilli, and a wedge of lime is ten. A piece of dark chocolate counts. Coffee counts. Green tea counts. Even a sprinkle of basil on top of a dish counts.

Once you start keeping track, thirty a week is not only achievable — it is surprisingly easy for anyone who cooks at home and uses a spice rack. For the person who lives on sandwiches, ready meals, and takeaway, it is almost impossible. That asymmetry is the most honest summary of what the modern microbiome literature is telling us about how food choices matter at the everyday level.

Why diversity, not dose

The why behind the thirty-plants finding is worth pausing over, because it tells us something about the nature of the gut ecosystem that the fibre story alone does not.

Each plant species contains a somewhat different mix of polysaccharides, polyphenols, proteins, and minor bioactive molecules. A carrot and a parsnip look similar on the plate but carry different classes of polysaccharide in their cell walls. Oats deliver beta-glucans that barley has less of; barley delivers arabinoxylans that oats have less of. A blueberry carries a different polyphenol profile from a blackberry, and both differ from a strawberry. Each of these slightly different substrates is best handled by a slightly different subset of gut microbes — different enzymes, different species, sometimes different whole metabolic guilds.

A person who eats the same four or five plants every day has, in effect, been running a one-crop farm in their colon. The microbes that can handle those four or five crops thrive. The microbes that require something else — a different polysaccharide, a different polyphenol, a different trace amount of some obscure glycan — cannot find work, and their populations dwindle. Over months and years, the ecosystem contracts around the narrow diet.

A person who eats thirty or forty plants a week is running something more like a mixed polyculture. Nothing is in huge abundance, but almost every microbial guild finds something to do. The ecosystem holds its diversity not because any one food is magical but because the *variety* creates niches for a wider range of inhabitants.

This is not speculation. A series of studies has now shown that increasing the *variety* of fibre types in a diet — rather than the total quantity — changes the microbiome and the metabolome in measurable ways. Ranaivo and colleagues, in a 2022 trial, showed that enriching a daily-consumed bread with a diverse rather than uniform mixture of fibre sources produced larger changes in microbial composition and metabolic markers in subjects at cardiometabolic risk, even when the total fibre dose was held constant [REF:ranaivo2022]. Diversity, in other words, is doing work that quantity alone cannot replicate.

The practical advice that falls out of this is refreshingly unlike most nutritional advice, because it does not require you to eliminate anything. It just asks you to add. When you plan a shopping list, reach for one new vegetable you have not bought before. When you cook, reach for one more herb than the recipe calls for. Keep seeds (sesame, pumpkin, sunflower, flax, chia) in the pantry

and sprinkle them on things. Eat the skins of fruits when they are edible; keep the leaves on the beetroot; buy whole spices and grind them yourself. Every species added is a new niche opened inside you.

Two things you can't see on the plate

There is a further reason to favour variety — two of them, in fact — that the thirty-plants number happens to hedge against even though the American Gut investigators were not looking for them. Both have to do with the quiet fact that the food on the modern plate is not quite what the food on the same plate was a few decades ago.

The first is what agronomists sometimes call the *dilution effect*. Donald Davis and colleagues at the University of Texas compared United States Department of Agriculture food-composition tables for forty-three garden crops between 1950 and 1999 and found statistically reliable declines in six of thirteen nutrients — protein, calcium, phosphorus, iron, riboflavin and ascorbic acid — on the order of six to thirty-eight per cent. The authors were careful in their own interpretation: the declines, they wrote, are “most easily explained by changes in cultivated varieties between 1950 and 1999, in which there may be trade-offs between yield and nutrient content” [REF:davis2004]. In other words, the same tomato or head of lettuce on today's supermarket shelf is very likely a different *cultivar* from the one your grandmother bought, bred over decades to pack more water, more starch, more sheer weight into each fruit, and proportionally less of several minerals and vitamins. The effect is not catastrophic and does not apply to every crop or every nutrient, but it is real enough, and consistent enough across foods, that it is difficult to dismiss. And unlike flashier concerns, it is quietly cumulative: a diet of five plants whose mineral density has slipped a third over two generations simply supplies less of those minerals, full stop.

Whether the drift matters for the microbiome is a reasonable next question, and the honest answer is: at the margin, in ways we are still learning to measure, probably yes. Gut bacteria have their own trace-element requirements, and the community does not merely follow the host's mineral status but actively shapes it. A 2021 review by Bielik and Kolisek catalogued the two-way traffic — dietary minerals shape which microbes thrive, and microbial metabolism in turn influences how much of those minerals the host can absorb [REF:bielik2021]. The surprising twist is that simply topping up a deficient nutrient in isolation does not automatically restore a healthy ecosystem, and can even make things worse. In a series of randomised trials in Kenyan infants, iron-fortified porridge reduced anaemia — a genuine benefit — but also shifted the colon toward more *Enterobacteriaceae* (including pathogenic *Escherichia coli*), reduced bifidobacteria, raised a biomarker of enterocyte damage, and increased treated respiratory infections. Adding a prebiotic (galacto-oligosaccharides) alongside the iron cancelled most of these adverse effects [REF:paganini2017]. A follow-up study in the same population showed that iron fortification even blunted the response to oral antibiotics and raised the risk of post-antibiotic diarrhoea [REF:paganini2018]. A broader review of excess iron in infants and young children reaches a consistent conclusion: iron is helpful when deficient, pro-inflammatory and dysbiotic when in surplus, and the margin between the two is narrower than many public-health programmes assume [REF:lonnerdal2017]. The lesson is not that fortification is bad — it is that a nutrient arriving on its own, without the food matrix it evolved inside, behaves differently from the same nutrient arriving embedded in a plant. The safer strategy — the one that captures whatever micronutrients our soils and cultivars still deliver, spreads the bet across many sources, and avoids accidental megadoses — is once again the thirty-plant rule.

The second hidden variable is residue. Of all the pesticides in routine modern use, the herbicide

glyphosate is the one the microbiome literature has scrutinised most closely, because its mechanism of action is not incidental to bacteria but directly relevant to them. Glyphosate kills plants by inhibiting a single enzyme, 5-enolpyruvylshikimate-3-phosphate synthase (EPSPS), which sits at the heart of the *shikimate pathway* — the route by which plants, fungi, and many bacteria build the aromatic amino acids tryptophan, tyrosine and phenylalanine. Humans lack the shikimate pathway entirely; we obtain these amino acids from food. This is the long-standing argument that glyphosate is safe for mammals. The catch, however, is that a large fraction of our gut bacteria *do* possess the pathway, and whether any particular species is sensitive or resistant depends on which *class* of EPSPS enzyme it carries. In honey bees, whose gut community is unusually tractable, Motta and colleagues showed that chronic exposure at environmentally realistic doses reshaped the community in a way that mapped precisely onto this distinction: species with the sensitive enzyme class declined, those with the resistant class held their ground, and the exposed bees became more susceptible to opportunistic infection [REF:motta2018]. The mechanism, in short, is not disputed.

Whether the doses that reach the human colon through food residue are sufficient to reproduce such an effect in a mammal is less settled, but the evidence is no longer nothing. In 2021, Mesnage and colleagues fed rats either glyphosate or the commercial Roundup formulation MON 52276 at doses bracketing the European regulatory limit, and combined shotgun metagenomics of the caecum with serum and caecal metabolomics. Both treatments caused shikimic acid and 3-dehydroshikimic acid — the substrates that accumulate when EPSPS is blocked — to build up in the gut contents, providing the first direct *in vivo* biochemical evidence that the shikimate pathway of the mammalian gut microbiome is in fact inhibited at doses within the legal range [REF:mesnage2021]. An earlier Ramazzini Institute study, using the United States Environmental Protection Agency's acceptable daily intake as the reference dose, found microbial and reproductive changes in the pups of rats whose mothers had been exposed during pregnancy [REF:mao2018]. The picture, however, is not one of straightforward devastation, and it would be dishonest to pretend otherwise. A 2022 biomonitoring study of British twins, which detected glyphosate in the urine of more than half of adult participants, actually found that higher excretion was weakly associated with *greater* microbial richness, alongside changes in several bacterial metabolites — a reminder that the effects are subtle, probably diet-dependent, and not a simple linear “more glyphosate, worse gut” story [REF:mesnage2022twins]. A 2023 review in *Gut Microbes* by Walsh and colleagues summarised the field as consistent with biologically plausible effects on community composition at realistic exposures, with genuine uncertainty about long-term and developmental outcomes [REF:walsh2023]. The appropriate response is neither panic nor dismissal. It is, once again, the ecological one: eat a wide enough variety of plants, from a wide enough variety of sources, that no single residue dominates the exposure — and, where the option is available and affordable, prefer sources (organic, small-farm, home-grown, seasonal) on which glyphosate is less likely to have been used in the first place.

Both of these stories — the dilution effect and the residue effect — are stories about things the modern eater cannot directly see or taste, and cannot realistically measure. Neither on its own is a crisis. But both are reasons why the answer to “how do I protect my microbiome from things I cannot even detect?” turns out to be the same as the answer to “how do I feed my microbiome well?” A diet of thirty different plant species a week is not a guarantee against either problem, but it is the simplest available hedge: it spreads the load across many sources, increases the chance that at least some of what you eat carries the minerals the tables promised, and dilutes the contribution of any single contaminated supply. Diversity, once again, is doing work that no single “superfood” and no supplement can replicate.

The redundancy that keeps the ecosystem upright

There is a second reason diversity matters that applies to ecosystems in general, including the one inside you: redundancy is resilience.

A forest with two dozen tree species is not only richer than a forest with three. It is also harder to destroy. If a pathogen arrives that attacks one species, the other twenty-three keep the forest standing. If a drought kills the shallow-rooted trees, the deep-rooted ones remain. The ecological jargon for this is “functional redundancy” — more than one species performs each key function, so losing any one member does not collapse the function.

The same holds for the gut. The production of butyrate, for example, is performed by a consortium of species — *Faecalibacterium prausnitzii*, *Roseburia intestinalis*, *Eubacterium rectale*, and many others. In a diverse microbiome, if one is temporarily suppressed (by an antibiotic, say, or a bout of illness), the others keep producing butyrate. In a monoculture gut where only one butyrate producer survives, losing that species means losing the function — and losing the SCFA means losing the colonocyte fuel and the Treg signal and the tight-junction integrity, all at once.

We cannot yet measure each reader’s redundancy directly. But we can predict with confidence that a gut supplied with thirty kinds of plant substrate a week will have more functional redundancy than one supplied with five. The thirty-plant rule is, in the end, a rule about building an ecosystem that can take a hit and stay upright.

20.5 The Mediterranean Pattern, Put to the Test

Why this one and not the others

The dietary patterns shelf of any bookshop is a chaos of competing claims. Keto, paleo, Whole30, DASH, Zone, Atkins, South Beach, carnivore, alkaline, blood-type, Gundry, Mediterranean, Nordic, Okinawan, MIND, intermittent fasting, flexitarian, and a dozen more — each with adherents, testimonials, and books that promise that this one, at last, is the answer.

Of all of them, the Mediterranean pattern is the only one that has been subjected, over decades, to the kind of testing that clinical medicine applies to a new drug: long-term randomised controlled trials with hard clinical endpoints, conducted in multiple countries, replicated across age groups, and scrutinised by independent review. It is also — and this is not a coincidence — the one that accidentally aligns most closely with the principles this chapter has been building up.

I say “accidentally” because the Mediterranean diet was not designed to feed a microbiome. It was observed in the mid-twentieth century by the American physiologist Ancel Keys, who noticed that people living on the Greek island of Crete and in certain regions of southern Italy had strikingly low rates of coronary heart disease despite eating enough calories and enough fat. Keys characterised what they were eating — plenty of vegetables, fruits, legumes, whole grains, nuts, olive oil, moderate amounts of fish, small amounts of meat and dairy (much of it fermented), modest amounts of wine with meals — and his characterisation became the template for the “Mediterranean diet” as it is now understood. The microbiome science that would eventually explain *why* it worked was still fifty years away.

What the pattern looks like on a plate is not exotic. It is mostly plants, in wide variety, lightly processed, eaten slowly, with olive oil as the dominant fat, and a small portion of animal protein

somewhere in the meal. It is Tuscan ribollita, Greek horiatiki, Provençal ratatouille, Lebanese fattoush, Turkish mezze. It is vegetables that taste like vegetables and bread that tastes like bread. A reader from almost any culture can find a version of it in their own food tradition.

NU-AGE: the big elderly trial

The most compelling recent evidence that the Mediterranean pattern works *through the microbiome* came from a study called NU-AGE, published in 2020 by Tarini Ghosh and Ian Jeffery at University College Cork working with collaborators across five European countries [REF:ghosh2020].

NU-AGE was a one-year randomised controlled trial in 612 elderly participants (aged 65 to 79) from Italy, France, the Netherlands, the United Kingdom, and Poland. The intervention group received personalised dietary counselling and tailored food deliveries to help them shift their eating towards a Mediterranean pattern. The control group continued their usual diet. At the end of the year, the researchers sequenced the gut microbiomes of both groups and measured a battery of biological markers associated with frailty — the progressive loss of muscle, strength, and metabolic resilience that defines healthy aging from unhealthy aging.

The result was that the participants who had shifted towards the Mediterranean pattern showed microbiome changes — increases in specific taxa including *Faecalibacterium prausnitzii*, *Roseburia*, and other SCFA producers; decreases in taxa associated with inflammation — that correlated, across the entire cohort, with reduced frailty markers and improved cognitive function. The changes were modest but real. More importantly, they were mediated by the microbiome: a statistical analysis showed that the effect of the diet on frailty was partly explained by the effect of the diet on specific microbial taxa and their metabolic outputs. This is what clinical epidemiologists call “mediation analysis,” and it is one of the cleanest ways to show that a mechanism is doing real work in a real-world intervention.

NU-AGE did not show a revolution. Elderly people put on a Mediterranean diet for a year did not suddenly become young. But they became measurably less frail than the controls, and the fingerprint of the microbiome change was visible in the data. For a dietary trial at this scale, in this age group, the result is unusual in its clarity.

Meslier: the shorter trial in a younger group

A second 2020 study, by Vanessa Meslier and colleagues working in Paris and Naples, asked a different question: in overweight and obese adults, does a Mediterranean diet intervention produce microbiome and metabolic changes independent of weight loss? [REF:meslier2020]

The study randomised 82 overweight or obese subjects to an 8-week isocaloric Mediterranean diet intervention or a habitual diet control. “Isocaloric” means the participants were not asked to eat less; they were asked to change *what* they ate without changing *how much*. This is an unusual design. Most diet trials bundle compositional change with caloric restriction, and it becomes impossible to tell whether the benefits come from the food composition or from simply eating fewer calories.

After eight weeks, the Mediterranean group had significant reductions in plasma cholesterol — roughly 10 milligrams per decilitre — with no change in body weight. The gut microbiome had shifted towards a more fibre-degrading, SCFA-producing composition. Plasma and faecal metabolomic profiles showed changes consistent with increased polyphenol and fibre intake. Because

weight had not changed, the benefits had to be coming from the change in what the subjects were eating rather than from weight loss.

This is an important finding because weight loss is the usual confounder in diet trials. Many dietary patterns produce temporary benefits that are actually benefits of transient caloric restriction, and the benefits vanish when weight is regained. The Meslier study carves out a cleaner signal: a diet change produced microbiome and metabolic improvements *without* a weight change, suggesting that the Mediterranean pattern has effects that are independent of its role in weight management. You do not have to lose weight to benefit from eating this way.

The pattern is a delivery system

Why does the Mediterranean pattern work so well as a microbiome intervention?

Because, under its Italian and Greek and Lebanese clothing, it is a delivery system for everything this chapter has so far argued matters. The fibre content is high, because it is a predominantly plant-based pattern with legumes, whole grains, vegetables, fruits, and nuts in almost every meal. The plant species count is high, because Mediterranean cuisines are built around herbs, spices, and mixed-vegetable preparations. The polyphenol content is high — we will come back to why in section 20.7 — because olive oil, red wine in moderation, coffee, tea, dark greens, and herbs are polyphenol powerhouses. The fermented-food content is meaningful, because the traditional Mediterranean diet includes yogurt, kefir, aged cheeses, olives, and in some regions sourdough breads. The saturated fat content is modest. The ultra-processed content is near zero in the traditional form, because the pattern predates the industrial food system that defines Western ultra-processed eating.

If you were designing a diet from first principles to feed a diverse, functional, SCFA-producing gut microbiome while also keeping the polyphenol and fermented-food contributions high, you would arrive somewhere very close to the Mediterranean pattern without ever having heard of it. It is the rare case in nutritional science where traditional wisdom and contemporary mechanism align so clearly that they seem to have been designed together.

One caveat is important. “Mediterranean diet” as sold in supermarkets, restaurants, and some commercial meal-kit services is often a costume rather than a substance. A “Mediterranean salad” drenched in commercial dressing, accompanied by white bread and a sugar-sweetened drink, is not the Mediterranean diet of Crete in 1960. The evidence base comes from the traditional pattern: mostly plants, in variety, with olive oil as the dominant fat, with small amounts of fish and even smaller amounts of red meat, with very little processed food, and with meals eaten slowly in social company. When this chapter recommends “a Mediterranean-style pattern,” it means the real thing, not the marketing label.

The last point — meals eaten slowly in social company — is not just cultural decoration. We will return to it in section 20.10, when we look at *when* you eat rather than *what* you eat. The Mediterranean pattern is not only a composition. It is also, in its traditional form, a rhythm. Both matter.

20.6 Fermented Foods: The Other Living Diet

A global habit, almost forgotten

Before refrigeration, almost every human culture preserved food by letting selected microbes grow on it. Milk became yogurt, kefir, and cheese. Cabbage became sauerkraut, kimchi, and a dozen regional relatives. Soy became miso, tempeh, natto, and soy sauce. Fish became the sauces of Southeast Asia and the garum of the Romans. Grains became sourdough bread and the starters behind every traditional beer. Tea became kombucha. Cassava became fufu. Somewhere between five and ten thousand years ago, humans learned that if you let the right microbes take over a food, the food would not only keep but also change — becoming more digestible, more flavourful, more nutritious, and, in many cases, better tolerated by people who could not drink plain milk or eat plain soy or chew plain cabbage leaf. Fermented foods, in this sense, are one of the oldest technologies in human history [REF:marco2017].

Then, over the course of the twentieth century, most of them quietly slipped off the Western plate. Refrigeration removed the practical necessity. Industrial production replaced live-culture yogurt with pasteurised “yogurt-style” desserts. Sauerkraut in most supermarkets is now pasteurised and shelf-stable — killed before it gets to you. Sourdough became the niche product of artisan bakeries while commercial bread switched to fast-acting baker’s yeast. The food tradition of microbial cooperation, maintained for millennia, thinned to a few pockets: natto in Japan, kimchi in Korea, kefir in parts of Eastern Europe, lambic beers in Belgium, and live yogurt in those dwindling corners where the word still means what it used to mean.

This matters for the microbiome in ways we have only recently begun to understand, and the evidence that has accumulated in the last few years is, in my view, the single most surprising development in the dietary microbiome literature. Not the fibre story — we knew fibre was important, and the details are refinements. The fermented-food story is closer to a discovery, and it came from one of the best-designed dietary trials the microbiome field has yet produced.

The Stanford trial

In 2021, a team led by Hannah Wastyk, Erica Sonnenburg, Justin Sonnenburg, and Christopher Gardner at Stanford published a paper in *Cell* that is, I think, a candidate for the most important dietary microbiome trial of the decade [REF:wastyk2021].

The design was clean in a way that trials of this kind rarely are. Thirty-six healthy adults were randomised into two arms. One arm was asked to progressively increase their fermented food intake over four weeks, working up from a baseline of near zero to a target of six servings of fermented foods per day, and then to hold that level for another six weeks — ten weeks total. The other arm was asked to progressively increase their fibre intake over the same period, working from a baseline of around 22 grams per day up to about 45 grams per day, and hold. Both interventions are, on paper, things we would expect to improve the microbiome, and the research team’s hypothesis was that both arms would show benefits, with the fibre arm probably doing more.

The team took stool samples, blood samples, and peripheral blood mononuclear cells at baseline, during the intervention, and after washout. They sequenced the microbiome in depth. They measured 93 circulating inflammatory markers. They profiled cytokine production by the participants’ own immune cells after ex vivo stimulation — a way of asking not just “are the inflammatory markers high?” but “how inflammatory is the immune system if you poke it?” In short, they measured not only the ecological change but also the physiological consequence.

Here is what they found, and it is worth slowing down to take it in.

The fermented-food arm showed a clear and significant increase in gut microbial diversity over the ten weeks. This is one of the few dietary interventions that has ever produced a measurable diversity increase in a randomised human trial. The fibre arm, at the same time, did *not* show an increase in diversity. The fibre arm did show other changes — shifts in the expression of microbial carbohydrate-degrading enzymes, among others — but the diversity needle did not move.

The immunology was even more striking. Nineteen of the 93 inflammatory markers the team measured went *down* significantly in the fermented-food arm. Markers associated with chronic low-grade inflammation — the kind implicated in metabolic disease, cardiovascular disease, and age-related frailty — dropped across the board. The immune cells of participants in the fermented-food arm were less reactive when stimulated, suggesting a reduction in baseline inflammatory tone. The fibre arm showed no such reduction. On the immune outcomes the trial was designed to detect, the fermented-food arm did what the researchers had predicted the fibre arm would do, and the fibre arm did not.

There are caveats. The trial was small — eighteen people per arm. It was only ten weeks. Both interventions are known to be beneficial, so the comparison is not “good versus bad” but “good versus very good.” And the finding that high-fibre intervention did not reduce inflammation in this trial contradicts earlier evidence from other trials, which forces us to be careful about generalising. One interpretation is that fibre needs longer than ten weeks to show its anti-inflammatory effect, particularly if the participant’s microbiome has been depleted by years of low-fibre eating and no longer contains the species that would do the fermenting. The Stanford team hinted at this: they pointed out that many of their fibre-arm participants showed evidence that their microbes simply could not handle the increased substrate, producing inflammatory fermentation byproducts rather than the expected shift to butyrate production. The machinery was not there. You can shovel oats into a gut whose fibre-fermenting species are missing, but the oats will not turn into butyrate. They will mostly just leave.

That interpretation is important for the practical section of the chapter, because it suggests that fermented foods may have a particular role for people whose microbiomes have been *depleted* — whose fibre-fermenting machinery is partly missing. Fermented foods deliver not only substrate but also live microbes, and at least some of those microbes, or their metabolic products, appear to do useful work.

What fermented foods actually do

The Stanford trial does not, on its own, tell us *how* fermented foods produce their effect. But we have a reasonable working picture, and it involves at least three mechanisms.

First, live microbes reach the distal gut. Not many of them. Not permanently. Most of the bacteria in a spoonful of yogurt are destroyed by stomach acid or outcompeted by the residents. But some arrive, some interact with the resident community, and some — particularly spore-formers and acid-tolerant species — persist long enough to do something useful. Taylor and colleagues, in a 2020 observational study, showed that habitual fermented-food consumers carry distinct microbial and metabolomic signatures compared with non-consumers, suggesting that the dietary habit leaves a measurable trace in the gut [REF:taylor2020].

Second, fermented foods deliver *microbial metabolites* that were produced during fermentation itself — short-chain peptides, bioactive amines, vitamins, exopolysaccharides, and compounds spe-

cific to the fermenting organism. When a yogurt culture ferments milk, it does not just thicken it; it enzymatically pre-digests some of the proteins, generates bioactive peptides from casein, produces lactic acid that alters the mineral availability of the product, and enriches the food with vitamins and precursors that the starting milk did not contain. You are not just eating dairy. You are eating the collected output of a microbial factory that worked on your behalf before you opened the container.

Third, fermentation *removes* components that would otherwise be problematic. Sourdough fermentation partially degrades the FODMAPs in wheat — the fermentable oligosaccharides that trigger bloating in many people — which is one reason that some individuals who feel bad after modern bread can tolerate traditional long-fermented sourdough without trouble. Kimchi and sauerkraut degrade some of the raw cabbage glucosinolates that cause gas. Tempeh and miso partially degrade phytates in soy, improving mineral bioavailability. Fermentation, in each case, is a kind of pre-processing that shifts the food from “raw plant” to “microbially prepared plant,” and the prepared version is often friendlier to the eater.

Getting six servings a day

Six servings a day sounds intimidating, but a serving, in the Stanford trial, was modest: a quarter-cup of kimchi or sauerkraut, half a cup of kefir or kombucha, a few ounces of yogurt, one or two pieces of aged cheese. The participants did not replace their normal meals with barrels of pickles. They added small amounts of fermented foods across the day: a dollop of yogurt at breakfast, a splash of kombucha mid-morning, a forkful of sauerkraut on the lunch plate, a few slices of a fermented cheese in the afternoon, a small portion of kimchi with dinner, a cup of kefir before bed. Six servings became achievable once the definition was right.

The practical advice, tiered by strength of evidence and by ease, comes out as follows. *If you do nothing else from this section, eat one serving of a real fermented food at most meals.* Yogurt with live cultures, kefir, traditional sauerkraut (refrigerated, not pasteurised, with visible bubbles or the note “contains live cultures” on the label), traditional kimchi, aged cheeses (which contain surviving microbial communities even though they are not typically described as fermented foods by lay readers), miso paste (added to soups after the soup is removed from heat to preserve the live cultures), tempeh, natto if you can find it and stomach it, kombucha (bearing in mind the sugar content of sweeter commercial varieties), and long-fermented sourdough bread from a bakery that actually ferments overnight rather than using commercial starter for flavour alone.

Avoid the supermarket impostors. “Pickles” produced by vinegar brining are not fermented — the vinegar is doing the preserving, not the microbes. “Yogurt” that lists “milk, cream, fruit, sugar, gelatin, pectin, natural flavours” without mentioning live cultures is almost certainly pasteurised after fermentation and contains no live organisms. Pasteurised sauerkraut from the canned-goods aisle is, biologically, cooked cabbage. Check the refrigerator section; read labels; look for visible bubbles or sediment; if in doubt, make your own. A batch of sauerkraut from one cabbage, one tablespoon of salt, and a jar takes about fifteen minutes to start and two weeks to finish. The resulting product is more expensive per jar than anything in the supermarket, because your time is valuable, but for roughly the cost of a cabbage you can build a supply of the real thing for a month.

The caveat paragraph

Before we move on, the obligatory caveat. Fermented foods are not a treatment for disease. The Wastyk trial was in healthy adults, and the outcomes were inflammatory markers rather than clin-

ical events. People with severely compromised immune systems — chemotherapy patients, transplant recipients, people on high-dose immunosuppressants — should talk to their clinicians before starting high-level fermented-food intake, because some live cultures can, rarely, cause problems in profoundly immunosuppressed hosts. People with histamine intolerance may find that aged cheeses and some fermented foods worsen their symptoms because fermentation produces biogenic amines. People on medications such as MAO inhibitors need to watch their tyramine intake, which is elevated in some aged cheeses and fermented soy products. These are real considerations for real people. The general population advice — most readers of this chapter — is that adding real fermented foods to the diet is one of the best-supported, lowest-risk, and most immediately effective things the microbiome literature currently points to.

20.7 Polyphenols: Why It Depends Whose Garden You Have

The chemistry of plant defence

Plants, unlike animals, cannot run away from what threatens them. They cannot dodge the insect that wants to chew them, the fungus that wants to rot them, or the sun that wants to oxidise their leaves. Instead, they fight chemically. Over four hundred million years of evolution, land plants have developed an extraordinary arsenal of secondary metabolites — molecules not directly involved in growth and reproduction, but deployed to deter enemies, attract allies, or protect against environmental stress. Many of these secondary metabolites are bitter, astringent, and vividly coloured, which is why the most chemically defended parts of plants often taste the most interesting and photograph the best.

A subset of these chemicals, called polyphenols, are defined by their structure: they carry multiple phenolic rings, chemical groups that can donate electrons and thereby neutralise reactive molecules. There are somewhere between eight and ten thousand known polyphenols, clustered into families: flavonoids (including the catechins of tea, the anthocyanins of berries, and the flavones of parsley and celery), phenolic acids (including the chlorogenic acid of coffee), stilbenes (including the resveratrol of red wine), lignans (in flaxseed and whole grains), and tannins (in wine, tea, and many fruit skins). Each family has its distinctive chemistry. What they share is that when you eat them, something surprising happens: most of them are not absorbed.

Estimates vary, but somewhere between 90 and 95 percent of dietary polyphenols pass through the small intestine and reach the colon intact or partly altered. They are too large to be readily absorbed across the intestinal wall, or too tightly bound to the plant's fibre matrix, or simply not in a form the enterocytes can transport. This fact used to be considered a problem — how can a compound that barely gets into the blood matter much for health? — until it became clear that the polyphenols are not meant for the enterocytes in the first place. They are meant for the microbes. And what the microbes do with them is where the story gets interesting [REF:cortes-martin2020].

Urolithin: the pomegranate story

Consider ellagitannins, a class of polyphenol concentrated in pomegranates, walnuts, raspberries, strawberries, and certain teas. When you eat a pomegranate, the ellagitannins arrive in your colon largely intact. There, some gut microbes — notably species in the *Gordonibacter* and *Ellagibacter* genera — possess the enzymes to break them down, releasing a smaller molecule called ellagic acid. Some microbes then process the ellagic acid further, through several intermediate steps, to produce

a compound called urolithin A. Urolithin A, unlike its parent molecules, is small enough and hydrophobic enough to be absorbed across the colonic wall. It enters the circulation and reaches other tissues. In cell and animal experiments, urolithin A has been shown to have effects on mitochondrial health, autophagy, and muscle function. Human trials, while early, suggest it may contribute to the benefits of polyphenol-rich diets in ageing-related muscle decline.

Here is the catch, and it is a big one. Not everybody makes urolithin A. Studies across multiple populations have found that people fall into three broad groups, which the research team of Francisco Tomás-Barberán in Murcia have called “metabotypes.” Metabotype A people convert ellagitannins efficiently into urolithin A. Metabotype B people produce mostly a different metabolite called urolithin B, which has weaker biological activity. Metabotype O people — and this is somewhere between 10 and 40 percent of the population, depending on the study — produce almost no urolithin at all. They simply do not carry the microbes that run the conversion, and their ellagitannins pass through them without yielding the useful downstream product [REF:tomas-barberan2017].

This means that the same handful of walnuts, eaten by three different people, produces three different systemic exposures to urolithin. For one person, the walnut is a delivery device for a potentially beneficial small molecule. For another, it delivers only the raw ingredient and the weaker derivative. For a third, it delivers nothing of the kind. The walnut is not different. The people are.

The equol story, and others

A similar pattern applies to soy isoflavones. Soy contains daidzein, one of the best-studied phytoestrogens. In some people, gut microbes convert daidzein into a molecule called equol, which has stronger oestrogenic activity than daidzein itself and appears to be the compound responsible for many of the effects attributed to soy — including, possibly, the reduced incidence of certain hormone-sensitive cancers in populations that eat soy from childhood. Only some people are equol producers. In Western populations, the proportion is often below 30 percent. In traditional Asian populations, the proportion tends to be higher, partly for genetic reasons but mostly, we think, because early-life dietary exposure and the associated microbial community favour the necessary converting species. If you are not an equol producer, the benefits attributed to soy in the epidemiological literature may not accrue to you in the same way.

The same principle applies, to varying degrees, to many other polyphenol classes. Flavan-3-ols from cocoa and green tea are processed by the microbes into phenyl- γ -valerolactones and other absorbable small molecules. Lignans from flaxseed are converted to enterolignans like enterodiol and enterolactone. Anthocyanins from berries are modified into a whole cascade of smaller phenolic metabolites. In each case, the ultimate systemic exposure to the biologically active product depends on whether the eater’s microbiome contains the necessary converting species, and the answer is frequently “partly” or “not much.”

What this means for you on Tuesday

It would be tempting to take this as a reason to give up on polyphenol-rich foods. After all, if you are not going to get the urolithin, why bother with the walnut? The answer is that the parent compounds, even without microbial conversion, carry useful effects of their own — direct antioxidant activity in the gut lumen, modification of the microbial community itself (several polyphenols preferentially feed beneficial taxa and suppress certain pathogens), and bioactivity from the smaller fraction that is absorbed intact. More importantly, the whole-food sources of polyphenols — berries, coloured vegetables, tea, coffee, dark chocolate, olives, olive oil, herbs and spices, red wine in

moderation, whole grains — are also, almost without exception, sources of fibre, vitamins, minerals, and plant diversity. Even if you are a metabotype O, a diet rich in polyphenolic foods is also a diet rich in everything else good.

The more sophisticated reading of the metabotype story is that it explains why population-level studies of polyphenol intake and disease sometimes produce inconsistent results. If the benefit depends on having a particular microbiome, then the average effect across a heterogeneous population will look smaller than the effect in the subset who can actually produce the active metabolite. This is a recurring theme in diet and microbiome research: the average hides the individuals, and some of the “inconsistencies” that have made nutritional epidemiology so frustrating over the decades may turn out to reflect real differences in who processes what.

The practical rule is simple. Eat polyphenol-rich foods in variety. Berries (fresh or frozen — frozen are often higher in polyphenols because they were picked ripe and preserved quickly). Tea, especially green tea. Coffee if you enjoy it. Dark chocolate (70% or higher, in small amounts, not as a substitute for actual food). Extra-virgin olive oil, which carries hydroxytyrosol and oleocanthal as well as fat. Olives. Herbs and spices — every culture’s traditional cuisine is a polyphenol delivery device disguised as flavouring. Red wine if it is part of your life, in modest amounts, with food. Nuts, particularly walnuts. Whole grains. Legumes, especially those with coloured seed coats. Do not be disappointed when headlines claim that this week’s superfood is “the answer”; there is no answer, only a rich and varied plate, and the microbes in your own particular gut will make of it what they can.

If, in five or ten years, commercial microbiome testing becomes cheap and reliable enough to tell you which polyphenol converters you carry, that information may become actionable. Today, it is not. Assume you are probably a partial converter for most classes, eat the foods anyway, and let your own particular ecosystem do its best.

20.8 The Other Side of the Plate: Protein, Fat, and the Metabolites You Don’t Want

TMAO: when your microbes turn dinner into a risk factor

In 2011, a group led by Stanley Hazen at the Cleveland Clinic published a paper in *Nature* that did something rare in cardiology: it added a new player to the list of molecules that predict heart disease, and the new player was not one of the usual suspects [REF:wang2011]. The molecule was trimethylamine-N-oxide, or TMAO, a small compound that had been known to chemists for a long time but had never been seriously considered as a cardiovascular risk factor until Hazen’s group showed that elevated blood levels of TMAO were strongly and independently predictive of heart attack, stroke, and cardiovascular death — even after adjusting for every traditional risk factor they could think of.

Where does TMAO come from? From your bacteria. Specifically, certain gut microbes break down dietary choline — abundant in eggs, red meat, poultry, and fish — into a gas called trimethylamine, or TMA. The TMA is absorbed into the portal vein, travels to the liver, and is oxidised by a host enzyme (FMO3) into TMAO, which then enters the systemic circulation. The TMAO is what seems to promote atherosclerosis, apparently by altering cholesterol handling in macrophages and contributing to foam-cell formation in the arterial wall.

In 2013, the Hazen group published a second paper, this one focused on a different precursor: L-carnitine, a molecule found in high concentrations in red meat [REF:koeth2013]. The mechanism was similar. Gut microbes convert carnitine to TMA; the liver oxidises it to TMAO; the TMAO contributes to atherosclerosis in mouse models and correlates with cardiovascular events in humans.

Here is the part of the Koeth study that made people sit up and pay attention. The researchers fed a steak (or, in the formal experiment, a standardised serving of beef) to a group of omnivores and measured their blood TMAO response. Levels rose. They then fed the same carnitine challenge to a group of long-term vegans and vegetarians. Levels barely moved. The vegans and vegetarians did not make TMAO from dietary carnitine, because they did not carry the bacterial populations that perform the conversion. Long-term avoidance of meat had, in effect, de-trained the microbial community for this particular substrate. Give it meat, and nothing happened. Give the omnivore meat, and the bacterial machinery that had been sustained by decades of meat-eating went to work.

This finding complicates the usual narrative about red meat and cardiovascular disease. The traditional interpretation — that meat is bad because of its saturated fat content — is partly true but incomplete. The Hazen-Koeth line of work suggests that a significant portion of the harm attributed to red meat is actually mediated by the microbial response to it, and that the microbial response itself depends on who you are and what you have been eating. The steak is not the only variable. The population of bacteria in your colon, shaped by years of previous dinners, is another.

It also suggests a form of partial reversibility. If you drop meat from your diet for long enough, the TMAO-producing populations shrink and the conversion slows. If you then eat meat occasionally, the conversion is partial. You do not have to become a lifelong vegan to benefit from reduced TMAO exposure; you may simply have to make meat something other than the centre of most meals. This is, again, what the Mediterranean pattern has been quietly telling us for sixty years: a small portion of fish or meat, surrounded by a large portion of plants, eaten a few times a week rather than most days, looks very different, in microbial terms, from a diet built around meat as the main event.

Saturated fat, taurine-bile, and the wadsworthia story

The second protein-and-fat story is more surprising, and it comes from a 2012 paper by Suzanne Devkota working in Eugene Chang's laboratory at Chicago [REF:devkota2012]. Devkota's study was in mice, but the mechanism it described is, as far as we can tell, fully applicable to the human bowel.

The experimental setup was this. Mice genetically engineered to be susceptible to inflammatory bowel disease (IL-10 knockout mice, missing a key anti-inflammatory signal) were placed on one of three diets: low fat, high fat with the fat coming from lard (mostly saturated), or high fat with the fat coming from safflower oil (mostly polyunsaturated). The team then measured the composition of the gut microbiota and the development of colitis.

On the low-fat diet, the susceptible mice mostly stayed healthy. On the polyunsaturated high-fat diet, they also mostly stayed healthy — so the problem was not fat per se. But on the saturated-fat diet, something striking happened. The composition of the mouse bile acid pool shifted. Normally mouse bile contains a mixture of glycine-conjugated and taurine-conjugated bile acids. On the high-saturated-fat diet, the taurine-conjugated fraction surged, because dealing with the higher fat load required more bile, and the taurine conjugates were preferentially produced. This was a subtle biochemical shift at the level of bile composition.

It was then amplified by the microbes. One bacterial species in particular, *Bilophila wadsworthia*,

thrives on taurine-derived sulphur. When the taurine-conjugated bile acids increased, *Bilophila wadsworthia* bloomed from a minor to a major component of the microbiota. Its metabolism produced hydrogen sulphide, a pro-inflammatory gas. The inflammatory burden on the colonic wall increased. And the genetically susceptible mice, in whom any additional inflammatory trigger could push the gut into overt disease, developed colitis.

The study is in mice with a specific genetic vulnerability, and we should not extrapolate it to every person who eats a cheeseburger. But the mechanism — saturated fat drives a bile acid shift that expands a sulphur-loving pathobiont that produces an inflammatory gas — is documented, biochemically coherent, and at least partially replicable in human studies that have since followed. *Bilophila wadsworthia* is present in the human gut and does expand on high-saturated-fat diets. Its sulphide production is real. Whether this produces frank colitis in most people is unclear; whether it contributes to the lower-grade inflammation associated with high-saturated-fat eating is plausible.

The speed of change

One more study in this cluster deserves a mention, because it tells us how fast dietary changes show up in the microbiome. In 2014, a team led by Lawrence David, then at Harvard, published a paper called “Diet rapidly and reproducibly alters the human gut microbiome” [REF:david2014]. The design was elegant. Ten healthy adults were fed, for five consecutive days, one of two extreme diets. The “plant-based” diet was essentially a vegan diet with high fibre and high complex carbohydrates. The “animal-based” diet contained meat, eggs, and cheeses at every meal, with almost no plant foods at all. Stool samples were collected before, during, and after each intervention.

Within 24 to 48 hours of starting either diet, the microbiome composition had shifted detectably. Within four days, the shifts were large. The animal-based diet promoted bile-tolerant organisms (*Bilophila*, among others) and suppressed Firmicutes that metabolise plant polysaccharides. The plant-based diet did the opposite. When the participants returned to their habitual diets, the changes reversed over a few days.

The take-home is both encouraging and sobering. Encouraging: your microbiome responds fast. A week of eating better is already starting to show results in the microbial community. Sobering: a week of eating badly has a similar speed of effect in the wrong direction, and if “a week of eating badly” becomes “a year of eating badly,” the cumulative shift can be substantial and, as we saw in the Sonnenburg extinction experiments, not fully reversible.

Enterotypes and the long term

There is one more layer to the picture, introduced by Gary Wu and collaborators at the University of Pennsylvania in 2011 [REF:wu2011]. They asked whether long-term dietary patterns — not what someone ate last week but what they had eaten for years — could be read off the gut microbiome. Their answer was yes, and they proposed a framework called “enterotypes”: broad categories of microbiome composition that cluster around different dominant genera.

The simplest version of the enterotype story, and the one that has survived the most scrutiny, is a spectrum between two extremes. At one end, people with a *Bacteroides*-dominated microbiome tend to have had long-term diets high in animal protein and animal fat. At the other end, people with a *Prevotella*-dominated microbiome tend to have had long-term diets high in fibre and plant-based carbohydrates. Short-term dietary change, as David showed, produces fast but partial shifts. Long-term dietary pattern produces deeper shifts that, in effect, select which ecosystem you carry.

Taken together, this cluster of studies — Hazen, Koeth, Devkota, David, Wu — tells a coherent story. What you eat today has fast effects on your gut microbes. What you eat across years shapes which microbes you carry, which metabolites they can make from your food, and how those metabolites act on the rest of your body. Animal foods are not the enemy. But a diet in which animal foods dominate over a long period trains a microbial community that produces more of the metabolites we don't want (TMAO, sulphides, secondary bile acids in inflammatory configurations) and less of the metabolites we do (butyrate and its SCFA cousins). The traditional Mediterranean pattern's recommendation of "meat a few times a week, plants at every meal" is, in hindsight, the right balance not because plants are virtuous and meat is sinful, but because this balance trains a microbial ecosystem whose metabolic outputs favour the host.

20.9 Ultra-Processed Food: When the Package Is the Problem

A category that didn't exist a generation ago

Walk through a modern supermarket and ask yourself a simple question: which of the items on these shelves would have been recognisable to your great-grandmother? The fruit and vegetable aisle, yes. The butcher and fish counter, mostly. The dry-goods aisle — flour, rice, dried beans — mostly. The dairy cabinet, with some caveats. Everything else — the coloured boxes, the plastic pouches, the brightly labelled tubs, the frozen ready meals, the "protein" bars, the "healthy" snacks, the "natural" beverages — contains products that did not exist in their current form before about 1970 and, for the most part, did not become a major fraction of the Western food supply until the 1980s and 1990s.

In 2009, a Brazilian epidemiologist named Carlos Monteiro proposed a classification system for foods based not on their nutrient content but on how much they had been industrially transformed. He called it the NOVA system. It has four categories. Group 1 is *unprocessed or minimally processed foods* — fruits, vegetables, grains, legumes, nuts, eggs, milk, meat, fish, spices, herbs. Group 2 is *processed culinary ingredients* — oils, butter, vinegar, salt, sugar, flour. Group 3 is *processed foods* — foods made by combining Group 1 and Group 2 items in recognisable ways, including traditional bread, cheese, canned vegetables, and cured meats. Group 4 is *ultra-processed foods*, and the definition is worth quoting, because the category is not simply "things with a lot of ingredients." Ultra-processed foods are, in Monteiro's formulation, industrial formulations that typically contain substances not found in home kitchens — protein isolates, modified starches, hydrogenated oils, high-fructose corn syrup, emulsifiers, stabilisers, artificial flavours, artificial colours, non-nutritive sweeteners — and that are designed to be hyper-palatable, shelf-stable, and profitable.

By this classification, ultra-processed foods now account for somewhere between 50 and 60 percent of total energy intake in the United States and the United Kingdom, and a substantial and rising fraction in most other industrialised countries. This is a dietary category that has moved, within a generation, from non-existent to dominant. And the evidence is accumulating, from multiple directions, that the microbiome is part of why it matters [REF:zinocker2018].

Emulsifiers: the wallpaper dissolver

Consider a product as innocent-sounding as commercial ice cream. Check the label. You will almost certainly find, somewhere in the ingredient list, one or more of: carboxymethylcellulose (also written CMC or E466), polysorbate 80 (E433), mono- and diglycerides, carrageenan, lecithin, or

xanthan gum. These are emulsifiers and stabilisers. Their job is to keep water and fat mixed smoothly, prevent ice crystals from forming, and give the product its characteristic mouthfeel. They are in ice cream, but they are also in commercial bread, salad dressings, plant milks, processed cheese, margarine, ready meals, protein shakes, and a substantial fraction of the processed food aisle. Most of them are regarded as safe in the sense that they do not cause acute toxicity at the doses used.

In 2015, Benoit Chassaing and Andrew Gewirtz at Georgia State published a paper in *Nature* that added a worrying footnote to that long-standing conclusion [REF:chassaing2015]. They fed mice CMC and polysorbate 80 at doses comparable to ordinary human exposure through the processed food supply — not experimental overdoses, but the kind of intake a person would get from a Western diet. They then measured microbiota composition, mucus layer integrity, inflammatory markers, and body weight.

The findings were, in a word, unsettling. The emulsifier-fed mice showed a thinning of the mucus layer — the same mucus layer we discussed in the fibre section, the layer that is supposed to separate the bacteria in the lumen from the epithelial cells of the gut wall. The bacteria, instead of living at a respectful distance, were now coming closer to the epithelium. Chassaing and Gewirtz called this “microbiota encroachment.” The microbiome composition shifted towards pro-inflammatory configurations. The mice developed low-grade inflammation, insulin resistance, and features of metabolic syndrome. In mice genetically predisposed to inflammatory bowel disease, frank colitis developed.

The reasonable scientific response was “this is in mice; do we see it in humans?” In 2022, Chassaing and colleagues published the answer: a randomised, controlled, double-blind human feeding trial in which healthy adults were given diets containing CMC at levels consistent with typical dietary exposure, or a matched control diet without CMC, for eleven days [REF:chassaing2022]. The sample was small — sixteen participants — but the experimental control was tight, and the outcomes were measured carefully.

The results confirmed, in humans, the direction of the mouse findings. Participants on the CMC diet showed reduced microbial diversity, alterations in the faecal metabolome, and — crucially — evidence of microbiota encroachment on the mucus layer in a subset of participants. Eleven days of an emulsifier that had been classified as safe for decades produced measurable changes in the way the gut microbes related to the gut wall.

This does not mean CMC is acutely dangerous. It does mean that a food additive which reaches us in small quantities through many products, day after day, year after year, is not as biologically inert as the regulatory history had assumed. And the list of emulsifiers that the Chassaing trial tested is only two molecules out of dozens that are routinely used in the ultra-processed food supply. We do not know, yet, what each of the others does. Polysorbate 80 behaves similarly to CMC in animal studies. Carrageenan has a longer-standing literature associating it with intestinal inflammation. Others have not been adequately studied.

The reasonable position, pending further data, is that dietary emulsifiers at Western food-supply levels are now a plausible contributor to the low-grade inflammation that characterises so much of contemporary chronic disease, and that reducing exposure to them is a sensible precaution even in the absence of definitive long-term outcome data. Practically, this means reading ingredient lists, favouring foods with short ingredient lists that contain only things you could buy at a market, and accepting that the creamy-mouthfeel ice cream of the supermarket freezer is not as trivial a treat as it looks on the label.

Non-nutritive sweeteners: the diet soda surprise

A parallel story, equally counterintuitive, has emerged around non-nutritive sweeteners. These are the compounds — saccharin, aspartame, sucralose, stevia glycosides, acesulfame potassium — that provide sweetness with few or no calories, and that have been marketed for decades as the healthy alternative to sugar. They were, for most of their history, considered microbiologically irrelevant on the theory that a compound which is not metabolised by the human body is also not metabolised by the gut microbiota.

In 2014, a group led by Eran Elinav and Eran Segal at the Weizmann Institute in Israel published a paper in *Nature* that turned this assumption upside down [REF:suez2014]. They showed that saccharin fed to mice at doses corresponding to the FDA's acceptable daily intake caused glucose intolerance — the precursor state to type 2 diabetes. They then showed that the mechanism was microbial: the saccharin altered the gut microbiota composition, and when faecal microbiota from saccharin-fed mice were transplanted into germ-free mice, the recipients developed glucose intolerance too. The microbial shift was the mediator. In a small pilot experiment in humans, a subset of volunteers who consumed saccharin for a week showed similar patterns — altered microbiota and impaired glucose responses — while another subset did not respond.

The 2014 paper was important but small and its human arm was limited. The definitive follow-up came in 2022, again from Suez, Elinav, and Segal, and was published in *Cell* [REF:suez2022]. This was a properly powered, double-blind, placebo-controlled human trial. One hundred and twenty healthy adults who had not been consuming non-nutritive sweeteners were randomised into six groups: saccharin, sucralose, aspartame, stevia, a placebo (glucose at a matched dose), or no supplement. The test sweeteners were given at doses well below the FDA's acceptable daily intakes. Outcomes included microbiome composition, metabolomic profiles, and glycemic response to oral glucose challenges.

The findings, taken carefully, are these. All four sweeteners altered the gut microbiome and the plasma metabolome, albeit in different directions. Two of them — saccharin and sucralose — also produced measurable impairment of glycemic responses in the participants who received them, with the impairment varying strongly between individuals. When faecal microbiota from the most responsive human participants were transplanted into germ-free mice, the mice recapitulated the glucose-intolerance phenotype — proof that the microbial changes were sufficient, on their own, to produce the glycemic effect.

Non-nutritive sweeteners, in other words, are not microbiologically neutral. They interact with the microbiome, the microbiome interacts with the host, and in some individuals the consequence is worse rather than better glucose handling. The person who drinks a diet cola in place of a sugar cola is not, on this evidence, getting the metabolic neutrality the marketing promises. Some people may be, some people are not, and we cannot currently tell in advance which category any given reader falls into.

The practical position for the general reader is cautious. If you are currently drinking multiple artificially sweetened beverages per day, it is worth considering whether the category is doing you the metabolic favour you had assumed. Reducing or eliminating them, substituting water, unsweetened tea and coffee, or a small amount of sugar, is a reasonable move. The same applies, more broadly, to any ultra-processed product whose “diet” or “light” credentials rest on sweetener substitution.

The Western diet as an ecological insult

The Zinöcker and Lindseth 2018 review pulled these strands — fibre depletion, emulsifier exposure, sweetener exposure, food-matrix disruption — into a single framework [REF:zinocker2018]. Their argument, in essence, is that the Western ultra-processed food system is best understood not as a specific nutrient problem but as a systemic ecological insult to the gut microbiome. Individual ingredients do individual damage. The combination — low MAC, constant emulsifier exposure, intermittent sweetener exposure, disrupted food matrices, displaced fermented foods, displaced plant diversity — acts on the ecosystem in multiple overlapping ways at the same time.

This is why single-nutrient approaches to improving dietary health — “eat less fat,” “eat more fibre,” “avoid sugar” — have historically produced modest and inconsistent results. The problem is not any one nutrient. The problem is the architecture. A diet built around ultra-processed products will be nutritionally inadequate in many overlapping ways, and tinkering with one nutrient while leaving the architecture intact produces small and often transient benefits. The alternative is not another single-nutrient rule. It is a structural shift: making the base of your diet minimally processed foods — Monteiro’s Groups 1, 2, and 3 — and letting Group 4 be a minor and occasional category rather than a dominant one.

I am aware that this sounds, in 2026, either obvious or impossible. Obvious, because it is what every sensible dietary writer for thirty years has been saying in different words. Impossible, because the convenient, affordable, and socially normal foods in most Western cities are predominantly ultra-processed, and asking a working family to “just cook from scratch” without acknowledging the time, money, and skill barriers that task involves is glib. I am not going to pretend those barriers do not exist. What I will say is that any move in the direction of minimally processed foods — even a partial one — is likely to pay returns, and that the returns will be disproportionately larger than equivalently sized changes within the ultra-processed category. Switching from one brand of packaged breakfast cereal to another does little. Replacing the breakfast cereal with oats and a handful of nuts does a lot. The difference is structural.

20.10 When You Eat: The Circadian Microbiome

The clock inside the clock

Until about a decade ago, most of us thought of the gut microbiome as a community that sat in the colon doing roughly the same thing around the clock. Eat a meal, digest it, excrete the remains, wait for the next meal. The composition of the community might shift with diet over weeks and months, but within a single day, we assumed, it was stable.

In 2014, Christoph Thaiss, then a graduate student in Eran Elinav’s laboratory at the Weizmann Institute, published a paper in *Cell* showing that this assumption was completely wrong [REF:thaiss2014]. Thaiss and colleagues sampled the gut microbiomes of mice at multiple time points across the 24-hour day and discovered that the composition oscillated. Specific bacterial taxa rose and fell in abundance over a daily rhythm. Specific microbial genes were transcribed more actively during some hours than others. The community was not static. It was cycling through daily phases that paralleled — and were driven by — the host’s feeding rhythm.

When the researchers disrupted the rhythm, by putting the mice on a jet-lag protocol (advancing the light cycle by 8 hours every three days), the microbial oscillations flattened and desynchron-

ised. The mice developed glucose intolerance and weight gain. When faecal microbiota from the jet-lagged mice were transplanted into germ-free mice on a normal schedule, the recipients also developed metabolic dysfunction — proof that the desynchronised microbiome was sufficient to cause the host metabolic phenotype. The team then studied two healthy human volunteers during actual transmeridian flights, and found evidence that the human gut microbiome was similarly desynchronised during jet lag and that the desynchrony correlated with the metabolic perturbations the volunteers experienced.

A follow-up paper from the same group in 2016 went further. Using mouse models, Thaïss and colleagues showed that the microbial oscillations were not just a passive consequence of feeding rhythms but an active driver of host physiology: the diurnal rhythms of microbial metabolite production shape the diurnal rhythms of gene expression in the liver and other tissues [REF:thaïss2016]. Your liver's daily schedule is not just set by your own clock. It is set, in part, by the metabolic output of the microbes in your gut, which in turn is set by when you eat.

This gave a microbial mechanism to a pattern that epidemiologists had been tracking for years. Shift workers have higher rates of metabolic syndrome, type 2 diabetes, cardiovascular disease, and some cancers than people who work regular hours. The usual explanation was that shift work disrupts sleep and stress systems. The Thaïss work added a new layer: shift work also disrupts the feeding rhythm, which disrupts the microbial rhythm, which disrupts the host metabolic rhythm. The damage is partly mediated by a microbial clock you cannot see.

Time-restricted eating

The obvious practical question is whether *when* you eat makes a measurable difference for people who are not shift workers. The growing body of human evidence suggests that it does, and the effect is large enough to take seriously.

Time-restricted eating is the practice of confining all daily food intake to a defined window — commonly eight, ten, or twelve hours — with the remaining portion of the 24-hour day spent not eating. It differs from the popular concept of “intermittent fasting” in that no calorie restriction is required; the person eats as much as they like within the window, but outside the window they drink only water, tea, or coffee without milk or sugar. In animal studies, time-restricted feeding produces improvements in glucose tolerance, lipid profiles, and body composition even when total calorie intake is held constant. In human studies, the evidence is more mixed but consistently points in the same direction.

A 2020 study by Zeb and colleagues randomised healthy adult men to a 16:8 time-restricted feeding pattern (8-hour eating window) or an ad libitum control for 25 days, and reported changes in gut microbiome composition, metabolic markers, and circadian rhythms in the restricted group that did not appear in the control [REF:zeb2020]. The effects were modest — this was a short trial in healthy participants, not a dramatic clinical intervention — but consistent with the broader picture.

More powerful evidence, at the population level, came from the French NutriNet-Santé cohort. In a 2023 paper, Palomar-Cros and colleagues analysed data from more than 100,000 adults whose eating times had been recorded prospectively [REF:palomar-cros2023]. They found that later first meal times and later last meal times were associated with higher cardiovascular disease risk, independent of total calorie intake and diet quality. Each additional hour of eating-window delay (eating later) was associated with meaningful increases in risk. Longer overnight fasts — essentially the space between dinner and breakfast — were associated with lower risk.

Put these threads together and you get the outline of a practical recommendation that, until recently, dietary advice did not contain. What matters is not only *what* you eat but also *when*. The simplest version of the recommendation is this: try to eat within a defined daily window, probably somewhere between 10 and 12 hours, and try not to eat in the hours immediately before bed. If you currently have your first coffee at 7 am and your last snack at 11 pm, you are eating across 16 hours, and your microbes — along with the circadian metabolism of your liver and your insulin-secreting cells — have no off-shift. Tightening the window to something like 8 am to 7 pm, or 9 am to 8 pm, gives the system a rest. For most people this is a free intervention: it does not cost anything, does not require changing what you eat, and can be adopted with relatively little disruption.

What the Mediterranean pattern had right all along

I want to return briefly to the Mediterranean pattern here, because one of its features that gets less attention than olive oil and tomatoes is the rhythm of its eating. In traditional Mediterranean cultures, meals were concentrated into three main events: a small breakfast, a significant mid-day meal eaten unhurriedly in company, and a light evening meal eaten relatively early. Snacking between meals was unusual. Eating late at night was unusual. The eating window was roughly 10 to 12 hours, not because anybody was practising time-restricted feeding but because that was the cultural pattern.

In other words, the pattern we now recommend for circadian reasons was already baked into the pattern we recommend for compositional reasons. This may be part of why long-term adherence to “the Mediterranean diet” as a lifestyle (not just a grocery list) produces results that tinkering with individual nutrients has historically not. The traditional pattern bundled the composition with the rhythm, and both parts may have been doing real work.

The person who wants the practical benefit without the full cultural immersion can take the simple advice: narrow the window, protect the overnight fast, avoid eating for two or three hours before sleep, and let the microbes have their daily rest.

20.11 Personalised Nutrition: Why “What Should I Eat?” Is the Wrong Question

The same banana, different people

In 2015, a team led by David Zeevi, Tal Korem, and Niv Zmora in the laboratories of Eran Segal and Eran Elinav at the Weizmann Institute published a study that I think should have changed the way popular nutrition is taught, and largely has not [REF:zeevi2015]. The study was called “Personalized Nutrition by Prediction of Glycemic Responses,” and it was, at the time, the largest and most detailed examination ever conducted of how real people respond to real food in their normal lives.

The design was this. Eight hundred healthy and pre-diabetic Israeli adults were recruited and equipped with continuous glucose monitors — the small skin-worn sensors that sample interstitial glucose every few minutes and produce a continuous record of blood glucose across days. For one week, the participants ate their normal diets, logged every meal in a smartphone app, and sent stool samples for microbiome sequencing. The researchers then standardised one meal per day — a specific bread, a specific breakfast cereal, a specific glucose drink — so that everybody was eating

the same thing at some point in the week. The team also collected anthropometric and clinical data: age, weight, height, blood tests, and so on. The total dataset included roughly 46,000 meals logged and sensor-tracked.

The analysis of that dataset produced a result so counter to the assumptions of public health nutrition that it is still underappreciated a decade later. When different people ate the *same* standardised meal — the same piece of bread, the same bowl of cereal, the same glucose drink — their blood glucose responses differed enormously. Person A's glucose might spike by 30 mg/dL after a slice of bread; person B's might spike by 10 mg/dL; person C's might spike by 90 mg/dL. These were not measurement errors. They were consistent, reproducible, within-person patterns that varied strikingly between people.

More interestingly, the variations were not predicted by any single simple variable. BMI, age, fasting glucose, and insulin resistance explained a little of the variance, but not most of it. The researchers then built a machine-learning model that incorporated all the measured variables — including the microbiome features — and asked whether the model could predict, for each individual, which foods would spike their glucose and which would not. It could. And the microbiome features were among the most important predictors. Particular gut bacterial taxa were associated with smoother glucose responses to particular foods. Other taxa were associated with sharper responses. The microbiome was carrying predictive information that the clinical variables were not.

What this means, in plain terms, is that two people can eat identical breakfasts and have entirely different metabolic experiences, and that the difference is partly a function of who their microbes are. “Bread causes a blood sugar spike” is true for some people and mostly false for others. “White rice is worse than brown rice” is a population-level statement that may not apply to any specific reader. The averaging has been hiding real variation, and the variation is not random — it is predictable, if you know enough about the individual.

PREDICT 1: the British twin study

The Weizmann work was done in a single population. It remained possible, in principle, that the striking individual variation was something about Israelis specifically, or about the particular measurement method, or about the chance characteristics of the sample. So in 2020, a large collaborative team — Tim Spector at King's College London, Andrew Chan at Harvard, Nicola Segata at Trento, and others — published the first results of a project called PREDICT 1, which was essentially the Weizmann design at larger scale and in a different population, with twins [REF:berry2020].

PREDICT 1 enrolled a thousand healthy adults, including hundreds of pairs of identical and non-identical twins, and subjected them to standardised meals under carefully controlled conditions, then tracked their free-living eating for two weeks. Continuous glucose, blood triglycerides, insulin, and other metabolic parameters were measured. The participants included twins because twin designs allow for partitioning of variance between genetic and non-genetic factors. If a trait varied more between non-identical twins than identical ones, you could quantify the genetic contribution.

The PREDICT 1 results replicated the Weizmann finding in a Western population and added precision to the partitioning. For postprandial glucose responses — the blood sugar curve after a meal — the contribution of genetics was small: perhaps 15 to 30 percent of the variance. Most of the individual variation was from non-genetic factors, including the microbiome, the person's current metabolic state, the time of day, the order in which components of the meal were eaten, the amount of sleep the person had had, and their habitual diet. For postprandial triglycerides, the pattern was

different — genetics contributed almost nothing to the variance, and the microbiome and lifestyle factors dominated almost completely.

A companion paper by Francesco Asnicar and colleagues, published in 2021 in *Nature Medicine*, specifically examined the microbiome arm of the PREDICT 1 data and reported that particular microbial taxa were associated with favourable and unfavourable cardiometabolic markers, and that habitual diet shaped microbiome composition in turn [REF:asnicar2021]. The team identified a “signature of 15 good microbes” that were associated with healthy plant-based diets and favourable blood markers, and a contrasting set of microbes associated with unhealthy processed diets and unfavourable markers. The key is that the relationships were mediated: diet shapes microbiome, microbiome shapes metabolism, metabolism shapes disease risk.

Where the evidence stands

Taken together, the Weizmann and PREDICT studies tell us that “one size fits all” nutrition is not a useful framework for individual-level advice. The same food produces different effects in different people, the differences are partly predictable from the microbiome, and the predictive models are good enough to be commercially interesting — several companies (including the ones descended from the Weizmann and PREDICT teams) now offer microbiome-based personalised nutrition services on a subscription basis.

But we should be careful here. A 2019 review by Kolodziejczyk, Zheng, and Elinav, published in *Nature Reviews Microbiology*, assessed the state of evidence for personalised nutrition and reached a balanced conclusion: the science is promising but the clinical use is still early [REF:kolodziejczyk2019]. The models predict group-level average responses better than individual-level variation. The tests are expensive. The actionable output — “eat more of X and less of Y” — is useful for some people in some contexts but does not, on current evidence, outperform general principles (more fibre, more plant diversity, less ultra-processed food) for most readers most of the time.

My own view, as a physician who watches how new technologies travel from research to clinic, is this. The scientific finding — that individual variation in response to food is real, substantial, and partly microbiome-mediated — is robust and important. The commercial translation — “send us a stool sample and we’ll tell you what to eat for £300” — is premature. If you have a specific clinical problem that is not responding to general advice (treatment-resistant type 2 diabetes, recurrent metabolic syndrome, stubborn weight gain despite reasonable eating), a personalised approach guided by continuous glucose monitoring and microbiome testing may add value. If you are a healthy reader trying to eat well, the general principles of this chapter — fibre, plant diversity, fermented foods, Mediterranean pattern, ultra-processed minimisation, time-restricted window — will do more, more cheaply, than any commercial test will.

The more interesting implication of the personalised nutrition literature, I think, is not that you should rush to buy a test. It is that you should stop trusting any single-food headline. When the Sunday newspaper tells you that “bananas cause blood sugar spikes” or “bananas are a perfect low-glycemic snack,” both statements are true and both statements are false, depending on which reader they are addressed to. The same goes for eggs, potatoes, white rice, oats, dairy, legumes, and almost every food that gets held up or held down in popular discussion. *You* are part of the equation, and the part of you that matters most for this particular question is sitting in your colon.

20.12 The Controversies, Met Head-On

I have avoided taking sides in the big contemporary food controversies until now, because the evidence-based case for the practical advice of this chapter does not depend on any of them. Fibre, plant diversity, fermented foods, Mediterranean pattern, minimising ultra-processed food, narrowing the eating window — all of that holds up regardless of where you stand on gluten, FODMAPs, lectins, or meat. But readers who have made it this far are entitled to a direct engagement with the positions they encounter on social media and in popular books, and this section is where I give them.

The structure of each of the following sub-sections is the same. I will state the claim as its adherents make it, describe what the best available evidence actually says, distinguish the kernel of truth from the exaggeration, and end with a practical recommendation. I will do this without sarcasm, because most of the people who hold these positions have arrived at them honestly — often because they felt unwell, tried something, felt better, and concluded that the thing they tried was the answer. That is a reasonable way to discover something. It is a less reasonable way to prove it.

Gluten

The claim made by the popular gluten-free movement is that wheat and other gluten-containing grains are, for a significant fraction of the population, a cause of bloating, fatigue, brain fog, joint pain, and a variety of systemic symptoms, and that the remedy is avoidance of all gluten-containing foods.

The evidence-based picture is more precise and more interesting. Celiac disease is a real, specific, autoimmune condition in which the body's own immune system attacks the lining of the small intestine in response to gluten exposure. It affects roughly 1 percent of most populations, it can be diagnosed reliably with blood tests and a small-bowel biopsy, and it requires strict lifelong gluten avoidance. If you have been formally diagnosed with celiac disease, the advice in this chapter about whole grains and fermented bread does not apply to you in the same form, and gluten-free alternatives are medically necessary rather than optional.

Non-celiac wheat sensitivity is also real, although how common it is depends on how it is defined. A proportion of people who do not have celiac disease but who report symptomatic improvement when they avoid wheat appear to have genuine sensitivity to *something* in wheat. The question that the research of the last decade has been trying to answer is: sensitivity to what, exactly?

The most carefully designed trial to address this question came from Jessica Biesiekierski, Peter Gibson, and colleagues at Monash University in 2013, and it tells a story that is not what most readers expect [REF:biesiekierski2013]. The Monash group recruited participants who self-identified as having non-celiac gluten sensitivity and whose symptoms had improved on a gluten-free diet. In a controlled crossover trial, participants were put on a low-FODMAP diet (that is, a diet low in fermentable oligosaccharides, disaccharides, monosaccharides, and polyols — we will come back to FODMAPs in a moment) and then randomised to receive either pure gluten, pure whey protein, or placebo, in a blinded fashion, added to an otherwise low-FODMAP background diet.

The result was striking. Participants' symptoms improved on the low-FODMAP diet. When they were then challenged with pure gluten, they reported no more symptoms than when challenged with whey protein or placebo. Gluten, when separated from the FODMAP content of wheat, did not reproduce the symptoms. FODMAP reduction, which was part of eliminating wheat, did. The trial suggested that a substantial fraction of the people who self-identify as "gluten sensitive" are

actually sensitive to the fermentable carbohydrates in wheat — particularly fructans — rather than to gluten itself.

The practical implication is important. If you feel better when you avoid wheat, you are not imagining it — something is probably going on. But the thing that is going on is more likely to be a FODMAP issue than a gluten issue, and the appropriate response is not necessarily to go gluten-free forever; it is to identify which fermentable carbohydrates you react to and consider long-fermented sourdough bread, which, because the yeasts and bacteria pre-digest much of the fructan content during the long rise, is often tolerated by people who cannot eat commercial bread.

A second piece of evidence is worth noting for anyone who has been tempted to adopt gluten avoidance as a general health measure rather than in response to specific symptoms. Benjamin Lebwohl and colleagues, in a 2017 *BMJ* paper, analysed long-term dietary and health data from over 100,000 American health professionals across more than two decades [REF:lebwohl2017]. They found that among people without celiac disease, long-term avoidance of gluten was associated with *higher*, not lower, cardiovascular disease risk, most plausibly because it was associated with reduced whole-grain intake. The gluten-free aisle of the supermarket is dominated by refined starches, added sugars, and industrial emulsifiers, and systematically replacing whole wheat with commercial gluten-free substitutes is likely to produce a nutritionally worse diet, not a better one.

So: if you have celiac disease, avoid gluten strictly. If you suspect non-celiac gluten sensitivity, do not assume gluten is the culprit; consider the FODMAP explanation, try long-fermented sourdough, and get formally assessed by a gastroenterologist or dietitian before committing to indefinite avoidance. If you have no symptoms and are considering “going gluten-free for health,” the evidence does not support it, and it may worsen your cardiovascular risk.

FODMAPs and the low-FODMAP diet

FODMAP stands for fermentable oligosaccharides, disaccharides, monosaccharides, and polyols — a collection of short-chain carbohydrates that are poorly absorbed in the small intestine, reach the colon, and are rapidly fermented by gut bacteria, producing gas and drawing water into the lumen. In people with irritable bowel syndrome (IBS), this gas and osmotic water shift can trigger bloating, pain, and altered bowel habit. The low-FODMAP diet, developed at Monash University, restricts dietary FODMAPs to a minimum and has become one of the best-supported dietary interventions for IBS symptom control.

The evidence for its efficacy in IBS is strong. Randomised trials, summarised in a 2017 review by Staudacher and Whelan, show symptom improvement in approximately 50 to 75 percent of IBS patients within four weeks of starting the strict elimination phase [REF:staudacher2017]. This is better than most pharmacological treatments for IBS, and the effect is generally rapid and clinically meaningful.

But the low-FODMAP diet has a problem from the microbiome point of view, and it is the same problem the fibre section of this chapter has been warning about. FODMAPs are fermentable substrates for the gut microbes — they are, in effect, a subset of microbiota-accessible carbohydrates. Stopping them reduces the substrate supply to the bacteria that ferment them, particularly *Bifidobacterium*, one of the most consistently beneficial genera in the human gut. A study by Halmos and colleagues in 2014 showed that a low-FODMAP diet for just three weeks significantly reduced faecal *Bifidobacterium* and total SCFA production [REF:halmos2014]. The short-term cost is real.

This is why the low-FODMAP diet is not meant to be a long-term way of eating, and when it is pre-

scribed by a knowledgeable dietitian it is explained as a three-phase process: (1) strict elimination for two to six weeks to achieve symptom control, (2) systematic reintroduction of one FODMAP class at a time to identify which ones trigger symptoms, and (3) long-term personalised restriction of the specific problematic FODMAPs only, with all tolerated FODMAPs returned to the diet. The third phase is the destination. The first phase is a diagnostic tool. People who get stuck in the first phase — who adopt the low-FODMAP elimination as a permanent lifestyle — are trading symptom control for accumulating microbiome impoverishment, and the trade may not be worth it long-term.

A related and interesting finding comes from Bennet and colleagues in 2018, who showed that baseline microbiome composition in IBS patients predicts who will respond to the low-FODMAP diet [REF:bennet2018]. Some IBS patients have microbial profiles that suggest they will respond well; others will not. This is another example of the personalisation theme — the same intervention, applied to different people, produces different results, and the difference is partly predictable from the microbiome.

Practical recommendation: if you have IBS and have not yet tried a properly supervised low-FODMAP diet, it is worth doing — under the guidance of a dietitian who understands the three-phase structure. If you have tried it and got stuck in the elimination phase, work with a dietitian to reintroduce systematically. If you do not have IBS, the low-FODMAP diet is not for you, and long-term FODMAP restriction in healthy people is likely to harm rather than help the gut microbiome.

Lectins

The claim, popularised by a series of bestselling books, is that lectins — a class of proteins found in many plants, particularly legumes, grains, nightshade vegetables, and some fruits — are a hidden cause of widespread chronic illness, and that a “lectin-free” diet is the path to weight loss, reduced inflammation, and improved health.

The evidence base for this claim is, to put it charitably, thin. Lectins are real proteins, and some of them — notably the ricin of castor beans — are genuinely toxic. Raw kidney beans contain a lectin called phytohaemagglutinin that causes acute gastrointestinal distress if eaten in quantity, which is why nobody eats raw kidney beans; a brief period of boiling completely denatures the lectin, and properly cooked beans are harmless. Raw wheat germ contains wheat germ agglutinin, which has been shown in laboratory experiments to interact with intestinal cells, but the quantity in normal bread and the effect of cooking reduce this to biological irrelevance.

The broader “lectin theory” extrapolates from these edge cases to the claim that lectins in ordinary, properly prepared, commonly consumed foods are causing harm. The epidemiological evidence runs powerfully against this claim. Legumes are one of the best-validated predictors of long-term health in every long-lived population that has been studied — the Blue Zone populations of Okinawa, Sardinia, and Nicoya all eat legumes in significant quantity and live exceptionally long lives. Whole grains are associated with lower all-cause mortality in every large cohort study that has measured them. Nightshade vegetables (tomatoes, peppers, eggplants, potatoes) are central to several of the healthiest traditional diets in the world. If lectins were causing the damage the theory posits, these associations would go the other way, and they do not.

From the microbiome point of view, the lectin theory is even less plausible, because the foods it demonises are, almost without exception, high in fibre and polyphenols and are precisely the foods we

have spent this chapter arguing are essential for microbiome health. Removing them would starve the very bacterial community that the rest of the popular “gut health” literature — including the same authors who promote lectin avoidance — claims to want to nurture. The position is internally inconsistent.

Practical recommendation: do not adopt lectin avoidance. Cook your beans properly, eat your whole grains, and let your tomatoes be. If you have a specific, formally diagnosed food intolerance, address that; do not generalise from it to an entire food category.

Carnivore

The most extreme of the current dietary movements is the carnivore diet: an eating pattern that eliminates all plant foods and consists exclusively of animal products, most commonly beef, eggs, butter, and small amounts of salt. Adherents report, often in dramatic terms, resolution of chronic symptoms — autoimmune disease, depression, inflammation, joint pain — on the elimination of plants.

I want to take these reports seriously, because dismissing them outright would be wrong. Some of the people who feel better on a carnivore diet have real symptoms that were not helped by conventional medicine, and their experience of improvement is real to them. The question is what the improvement actually represents.

The most plausible explanation — and this is a hypothesis, not an established fact — is that for a subset of these individuals, what appears to be an improvement on “removing plants” is actually an improvement on removing a specific subset of plant compounds to which they were reacting. If that subset included the fermentable carbohydrates that cause IBS, the oxalates that precipitate in some people’s joints or kidneys, the histamine-liberating foods that provoke mast-cell reactions, the nightshades some people with autoimmune conditions are sensitive to, and perhaps the phytate-chelated minerals that are poorly absorbed in some gut profiles — then removing all of them at once by removing all plants could produce the improvement without the improvement being specifically due to “animal foods.” The same improvement, in those individuals, might well be achievable by a carefully designed plant-inclusive diet that happens to avoid their specific triggers.

The problem with the carnivore diet as a long-term strategy is that we have almost no long-term data. The evidence base consists of testimonials, short-term observational reports, and theoretical argument. What we know about gut microbiome biology suggests that a strict and sustained plant-free diet will produce a microbiome with markedly reduced diversity, near-absent fibre-fermenting capacity, shrunk SCFA production, and correspondingly reduced butyrate fuel for colonocytes. We have already discussed what the David 2014 diet-switch experiment showed in humans: dramatic shifts within days towards animal-diet-typical microbial profiles, including expansion of bile-tolerant and sulphide-producing organisms [REF:david2014]. We have already discussed what the TMAO literature shows about heavy meat intake and cardiovascular risk [REF:koeth2013]. We have already discussed the Sonnenburg extinction experiments and what low-MAC diets do to diversity across generations [REF:sonnenburg2016]. None of these are reassuring about a prolonged carnivore pattern.

My practical recommendation is this. If you feel unwell on your current diet and have not been able to get good answers from conventional medicine, the carnivore experiment is not unreasonable as a short-term diagnostic tool — two or three weeks of strict elimination, followed by systematic reintroduction of plant categories one at a time, might help you identify what you actually react

to. This is, in essence, the same diagnostic logic as the low-FODMAP diet, just starting from a different end. What I would not recommend is the indefinite maintenance of the strict elimination phase as a long-term way of eating, on testimonial evidence alone, against a body of mechanistic and epidemiological evidence pointing the other way. The upside — that you personally will be an exception to everything we know about long-term nutrition and microbiome biology — is possible but uncertain. The downside — that you will de-train a microbial community across years in ways that cannot be easily undone — is documented.

A note on certainty

I have stated positions in this section more directly than I usually would in a book like this, because I think the reader is entitled to them. I want to be equally direct that in each of these cases, the evidence base has gaps and the scientific conversation continues. If new high-quality trials change the picture — if a properly designed long-term carnivore trial showed benefits we did not predict, or if the FODMAP explanation for non-celiac sensitivity turned out to be only part of the story — I would update my position, and a later edition of this book would say something different. What I have tried to give you is the current best reading of the best available evidence, not a pronouncement from on high. Take it as such, and watch the literature.

20.13 What This Means for Tuesday Dinner

Here is the practical summary of the chapter, tiered by the strength of the evidence, so that a reader who has read only this section has the useful distillation.

The strongest evidence: things to do

Eat more fibre. Work up to 30 grams or more per day. Choose a variety of sources — whole grains, legumes, vegetables, fruits, nuts, seeds — rather than relying on a single fibre type. If you are currently at 12 to 15 grams per day (which is most Western adults), doubling your intake over a few weeks is one of the single most evidence-based changes you can make.

Aim for thirty different plant species per week. Count everything: vegetables, fruits, whole grains, legumes, nuts, seeds, herbs, spices, coffee, tea, dark chocolate. Write them down for a week if you find it hard to believe you are reaching or not reaching the target. The number itself is less important than the habit of variety.

Eat real fermented foods most days. A small serving — a spoonful of yogurt with live cultures, a forkful of kimchi or sauerkraut from the refrigerated section, a cup of kefir, a piece of aged cheese, a small glass of kombucha — at most meals. Read labels; avoid the pasteurised impostors. The Stanford evidence suggests this may be the single intervention most likely to reduce inflammatory tone in the healthy adult.

Adopt a Mediterranean-style pattern. Plants at the centre of the plate, olive oil as the dominant fat, small portions of fish and meat a few times a week, legumes and whole grains at most meals, herbs and spices in abundance, fruit for dessert, water and unsweetened drinks as the default beverage. This is not a diet in the sense of a temporary restriction; it is a cuisine, and there are a hundred versions of it from a hundred cultures.

Eat within a defined daily window. Somewhere between 10 and 12 hours is a reasonable target for most people. Protect the overnight fast; avoid eating in the two or three hours before sleep.

Strong evidence: things to reduce

Reduce ultra-processed foods. The category includes almost everything in brightly coloured packaging with long ingredient lists and names you cannot pronounce. Perfect elimination is not required and often not practical; steady reduction is. Aim to get the majority of your calories from foods recognisable as foods. A useful heuristic: if it was not food a hundred years ago, be suspicious.

Reduce dietary emulsifiers. Read labels for carboxymethylcellulose (E466), polysorbate 80 (E433), carrageenan, and mono- and diglycerides. These are concentrated in processed ice creams, commercial breads, processed cheeses, plant milks, salad dressings, and ready meals. You cannot avoid them entirely and probably do not need to, but reducing daily exposure is sensible.

Reduce non-nutritive sweeteners. Particularly saccharin and sucralose, the two with the clearest evidence of microbiome-mediated effects in the Suez 2022 human trial. Aspartame and stevia have less conclusive evidence but also produced microbiome changes in the same study and should be viewed with modest caution. Water, unsweetened tea, unsweetened coffee, and small amounts of real sugar are all better choices than constant diet-sweetened drinks.

Reduce red and processed meat to a few times a week at most. Not because meat is poison but because heavy long-term intake trains a microbial community that produces more TMAO and secondary bile acid metabolites than most of us want. A small serving of good meat, a few times a week, as part of a plant-heavy meal, is consistent with everything this chapter has argued for.

Promising but not yet prescription-grade

Personalised microbiome-based meal planning. The science is interesting and improving. The commercial services are expensive and the output is not yet clearly better than general principles for most people. Wait, for now, unless you have a specific clinical problem that general advice has not solved.

Polyphenol supplementation. Eat the whole foods; supplementation with isolated polyphenols has produced inconsistent results in trials, and the metabotype story makes it clear that the systemic effect depends on who is eating.

Extreme time-restricted feeding windows (shorter than 8 hours). The evidence for a 10- to 12-hour window is reasonable. The evidence for shorter windows is less clear, and very short windows (4 to 6 hours) may cause problems for some people, particularly women, shift workers, and people with blood sugar regulation issues. Start gentle.

Not supported by evidence

Gluten-free eating for people without celiac disease or diagnosed wheat allergy. Does not help; may harm.

Lectin avoidance. Eliminates some of the most strongly evidence-based foods on the planet on the basis of a theory not supported by human epidemiology.

Strict long-term carnivore eating. Short-term diagnostic use may have a role. Long-term maintenance is not supported by existing evidence and conflicts with much of what we know about microbiome biology.

Fad diets generally. If a diet is named after its inventor, promises dramatic results in weeks, demands a book purchase to understand fully, and excludes one or more major food groups, be cautious. The evidence base for any of these, compared with the traditional Mediterranean pattern or any whole-foods plant-heavy approach, is comparatively thin.

The closing image

Dinner is a daily act of ecology. When you choose what to put on your plate, you are not only feeding yourself. You are feeding a forest — a living, breathing, metabolising community of trillions of organisms whose daily output of small molecules shapes how your own cells function, how your immune system responds, how your gut wall holds together, and in the long run how well you age. Most of them you will never see and never think about. But they will process every meal you eat, and what they make of it depends, every day, on what you choose to give them.

The garden metaphor I began this chapter with has one last application. A gardener who tends a piece of ground for decades does not produce miraculous results through any single intervention. She produces them through the accumulated consequence of small, consistent choices across many seasons. The right compost. The right watering. The right crop rotation. The right attention to the slow signals of the soil. None of it is dramatic. All of it matters. And at the end of twenty years, one gardener has a rich, productive, resilient plot and another has depleted earth, and the difference is the sum of daily decisions that neither of them noticed mattering at the time.

Your gut is the ground you have to work with. Feed it well, and its inhabitants will pay you back across a lifetime.

20.14 Looking Ahead

Diet is the single most powerful daily lever most of us have over our microbiome. But it is not the only lever, and it sits inside a much larger web of environmental and lifestyle factors that shape the community of organisms we carry. The air you breathe, the water you drink, the dirt your children play in, the pets in your home, the people you live with, the medications you take, the sleep you get, the stress you carry, the exercise you do or do not do — all of them leave fingerprints on the ecosystem inside you. Chapter 21 widens the lens to look at those other inputs, and asks how the rest of the way we live — beyond what we put on the plate — is shaping the microbial communities that are, by now, unmistakably part of who we are.

Key References

Asnicar F, Berry SE, Valdes AM, et al. (2021). Microbiome connections with host metabolism and habitual diet from 1,098 deeply phenotyped individuals. *Nature Medicine* 27(2):321-332. doi:10.1038/s41591-020-01183-8. [REF:asnicar2021]

- Bennet SMP, Böhn L, Störsrud S, et al. (2018). Multivariate modelling of faecal bacterial profiles and food intake in irritable bowel syndrome. *Gut* 67(5):872-881. doi:10.1136/gutjnl-2016-313128. [REF:bennet2018]
- Berry SE, Valdes AM, Drew DA, et al. (2020). Human postprandial responses to food and potential for precision nutrition. *Nature Medicine* 26(6):964-973. doi:10.1038/s41591-020-0934-0. [REF:berry2020]
- Bielik V, Kolisek M (2021). Bioaccessibility and bioavailability of minerals in relation to a healthy gut microbiome. *International Journal of Molecular Sciences* 22(13):6803. doi:10.3390/ijms22136803. [REF:bielik2021]
- Biesiekierski JR, Peters SL, Newnham ED, Rosella O, Muir JG, Gibson PR (2013). No effects of gluten in patients with self-reported non-celiac gluten sensitivity after dietary reduction of fermentable, poorly absorbed, short-chain carbohydrates. *Gastroenterology* 145(2):320-328.e3. doi:10.1053/j.gastro.2013.04.051. [REF:biesiekierski2013]
- Chassaing B, Koren O, Goodrich JK, et al. (2015). Dietary emulsifiers impact the mouse gut microbiota promoting colitis and metabolic syndrome. *Nature* 519(7541):92-96. doi:10.1038/nature14232. [REF:chassaing2015]
- Chassaing B, Compher C, Bonhomme B, et al. (2022). Randomized controlled-feeding study of dietary emulsifier carboxymethylcellulose reveals detrimental impacts on the gut microbiota and metabolome. *Gastroenterology* 162(3):743-756. doi:10.1053/j.gastro.2021.11.006. [REF:chassaing2022]
- Cortés-Martín A, Selma MV, Tomás-Barberán FA, González-Sarriás A, Espín JC (2020). Where to look into the puzzle of polyphenols and health? The postbiotics and gut microbiota associated with human metabolotypes. *Molecular Nutrition & Food Research* 64(9):e1900952. doi:10.1002/mnfr.201900952. [REF:cortes-martin2020]
- David LA, Maurice CF, Carmody RN, et al. (2014). Diet rapidly and reproducibly alters the human gut microbiome. *Nature* 505(7484):559-563. doi:10.1038/nature12820. [REF:david2014]
- Davis DR, Epp MD, Riordan HD (2004). Changes in USDA food composition data for 43 garden crops, 1950 to 1999. *Journal of the American College of Nutrition* 23(6):669-682. doi:10.1080/07315724.2004.10719409. [REF:davis2004]
- Deehan EC, Mocanu V, Madsen KL (2024). Effects of dietary fibre on metabolic health and obesity. *Nature Reviews Gastroenterology & Hepatology* 21(5):301-318. doi:10.1038/s41575-023-00891-z. [REF:deehan2024]
- Desai MS, Seekatz AM, Koropatkin NM, et al. (2016). A dietary fiber-deprived gut microbiota degrades the colonic mucus barrier and enhances pathogen susceptibility. *Cell* 167(5):1339-1353.e21. doi:10.1016/j.cell.2016.10.043. [REF:desai2016]
- Devkota S, Wang Y, Musch MW, et al. (2012). Dietary-fat-induced taurocholic acid promotes pathobiont expansion and colitis in *Il10*^{-/-} mice. *Nature* 487(7405):104-108. doi:10.1038/nature11225. [REF:devkota2012]
- Furusawa Y, Obata Y, Fukuda S, et al. (2013). Commensal microbe-derived butyrate induces the differentiation of colonic regulatory T cells. *Nature* 504(7480):446-450. doi:10.1038/nature12721. [REF:furusawa2013]

- Ghosh TS, Rampelli S, Jeffery IB, et al. (2020). Mediterranean diet intervention alters the gut microbiome in older people reducing the risk of frailty: results from the NU-AGE 1-year dietary intervention across five European countries. *Gut* 69(7):1218-1228. doi:10.1136/gutjnl-2019-319654. [REF:ghosh2020]
- Halmos EP, Christophersen CT, Bird AR, Shepherd SJ, Gibson PR, Muir JG (2014). Diets that differ in their FODMAP content alter the colonic luminal microenvironment. *Gut* 63(1):93-100. doi:10.1136/gutjnl-2014-307264. [REF:halmos2014]
- Koeth RA, Wang Z, Levison BS, et al. (2013). Intestinal microbiota metabolism of L-carnitine, a nutrient in red meat, promotes atherosclerosis. *Nature Medicine* 19(5):576-585. doi:10.1038/nm.3145. [REF:koeth2013]
- Koh A, De Vadder F, Kovatcheva-Datchary P, Bäckhed F (2016). From dietary fiber to host physiology: short-chain fatty acids as key bacterial metabolites. *Cell* 165(6):1332-1345. doi:10.1016/j.cell.2016.05.041. [REF:koh2016]
- Kolodziejczyk AA, Zheng D, Elinav E (2019). Diet-microbiota interactions and personalized nutrition. *Nature Reviews Microbiology* 17(12):742-753. doi:10.1038/s41579-019-0256-8. [REF:kolodziejczyk2019]
- Lebwohl B, Cao Y, Zong G, et al. (2017). Long term gluten consumption in adults without celiac disease and risk of coronary heart disease: prospective cohort study. *BMJ* 357:j1892. doi:10.1136/bmj.j1892. [REF:lebwohl2017]
- Lönnerdal B (2017). Excess iron intake as a factor in growth, infections, and development of infants and young children. *American Journal of Clinical Nutrition* 106(Suppl 6):1681S-1687S. doi:10.3945/ajcn.117.156042. [REF:lonnerdal2017]
- Mao Q, Manservigi F, Panzacchi S, et al. (2018). The Ramazzini Institute 13-week pilot study on glyphosate and Roundup administered at human-equivalent dose to Sprague Dawley rats: effects on the microbiome. *Environmental Health* 17(1):50. doi:10.1186/s12940-018-0394-x. [REF:mao2018]
- Marco ML, Heeney D, Binda S, et al. (2017). Health benefits of fermented foods: microbiota and beyond. *Current Opinion in Biotechnology* 44:94-102. doi:10.1016/j.copbio.2016.11.010. [REF:marco2017]
- McDonald D, Hyde E, Debelius JW, et al.; American Gut Consortium (2018). American Gut: an open platform for citizen science microbiome research. *mSystems* 3(3):e00031-18. doi:10.1128/mSystems.00031-18. [REF:mcdonald2018]
- Meslier V, Laiola M, Roager HM, et al. (2020). Mediterranean diet intervention in overweight and obese subjects lowers plasma cholesterol and causes changes in the gut microbiome and metabolome independently of energy intake. *Gut* 69(7):1258-1268. doi:10.1136/gutjnl-2019-320438. [REF:meslier2020]
- Mesnage R, Teixeira M, Mandrioli D, et al. (2021). Use of shotgun metagenomics and metabolomics to evaluate the impact of glyphosate or Roundup MON 52276 on the gut microbiota and serum metabolome of Sprague-Dawley rats. *Environmental Health Perspectives* 129(1):17005. doi:10.1289/EHP6990. [REF:mesnage2021]
- Mesnage R, Bowyer RCE, El Balkhi S, et al. (2022). Impacts of dietary exposure to pesticides on faecal microbiome metabolism in adult twins. *Environmental Health* 21(1):46.

doi:10.1186/s12940-022-00860-0. [REF:mesnage2022twins]

Motta EVS, Raymann K, Moran NA (2018). Glyphosate perturbs the gut microbiota of honey bees. *Proceedings of the National Academy of Sciences USA* 115(41):10305-10310. doi:10.1073/pnas.1803880115. [REF:motta2018]

Paganini D, Uyoga MA, Kortman GAM, et al. (2017). Prebiotic galacto-oligosaccharides mitigate the adverse effects of iron fortification on the gut microbiome: a randomised controlled study in Kenyan infants. *Gut* 66(11):1956-1967. doi:10.1136/gutjnl-2017-314418. [REF:paganini2017]

Paganini D, Uyoga MA, Kortman GAM, et al. (2018). Iron-containing micronutrient powders modify the effect of oral antibiotics on the infant gut microbiome and increase post-antibiotic diarrhoea risk: a controlled study in Kenya. *Gut* 68(4):645-653. doi:10.1136/gutjnl-2018-317399. [REF:paganini2018]

Palomar-Cros A, Andreeva VA, Fezeu LK, et al. (2023). Dietary circadian rhythms and cardiovascular disease risk in the prospective NutriNet-Sant  cohort. *Nature Communications* 14(1):7899. doi:10.1038/s41467-023-43444-3. [REF:palomar-cros2023]

Ranaivo H, Thirion F, B ra-Maillet C, et al. (2022). Increasing the diversity of dietary fibers in a daily-consumed bread modifies gut microbiota and metabolic profile in subjects at cardiometabolic risk. *Gut Microbes* 14(1):2044722. doi:10.1080/19490976.2022.2044722. [REF:ranaivo2022]

Reynolds A, Mann J, Cummings J, Winter N, Mete E, Te Morenga L (2019). Carbohydrate quality and human health: a series of systematic reviews and meta-analyses. *Lancet* 393(10170):434-445. doi:10.1016/S0140-6736(18)31809-9. [REF:reynolds2019]

Schnorr SL, Candela M, Rampelli S, et al. (2014). Gut microbiome of the Hadza hunter-gatherers. *Nature Communications* 5:3654. doi:10.1038/ncomms4654. [REF:schnorr2014]

Smith PM, Howitt MR, Panikov N, et al. (2013). The microbial metabolites, short-chain fatty acids, regulate colonic Treg cell homeostasis. *Science* 341(6145):569-573. doi:10.1126/science.1241165. [REF:smith2013]

Smits SA, Leach J, Sonnenburg ED, et al. (2017). Seasonal cycling in the gut microbiome of the Hadza hunter-gatherers of Tanzania. *Science* 357(6353):802-806. doi:10.1126/science.aan4834. [REF:smits2017]

Sonnenburg ED, Sonnenburg JL (2014). Starving our microbial self: the deleterious consequences of a diet deficient in microbiota-accessible carbohydrates. *Cell Metabolism* 20(5):779-786. doi:10.1016/j.cmet.2014.07.003. [REF:sonnenburg2014]

Sonnenburg ED, Smits SA, Tikhonov M, Higginbottom SK, Wingreen NS, Sonnenburg JL (2016). Diet-induced extinctions in the gut microbiota compound over generations. *Nature* 529(7585):212-215. doi:10.1038/nature16504. [REF:sonnenburg2016]

Staudacher HM, Whelan K (2017). The low FODMAP diet: recent advances in understanding its mechanisms and efficacy in IBS. *Gut* 66(8):1517-1527. doi:10.1136/gutjnl-2017-313750. [REF:staudacher2017]

Suez J, Korem T, Zeevi D, et al. (2014). Artificial sweeteners induce glucose intolerance by altering the gut microbiota. *Nature* 514(7521):181-186. doi:10.1038/nature13793. [REF:suez2014]

Suez J, Cohen Y, Vald s-Mas R, et al. (2022). Personalized microbiome-driven effects of non-nutritive sweeteners on human glucose tolerance. *Cell* 185(18):3307-3328.e19. doi:10.1016/j.cell.2022.07.016.

[REF:suez2022]

Taylor BC, Lejzerowicz F, Poirel M, et al. (2020). Consumption of fermented foods is associated with systematic differences in the gut microbiome and metabolome. *mSystems* 5(2):e00901-19. doi:10.1128/mSystems.00901-19. [REF:taylor2020]

Thaiss CA, Zeevi D, Levy M, et al. (2014). Transkingdom control of microbiota diurnal oscillations promotes metabolic homeostasis. *Cell* 159(3):514-529. doi:10.1016/j.cell.2014.09.048. [REF:thaiss2014]

Thaiss CA, Levy M, Korem T, et al. (2016). Microbiota diurnal rhythmicity programs host transcriptome oscillations. *Cell* 167(6):1495-1510.e12. doi:10.1016/j.cell.2016.11.003. [REF:thaiss2016]

Tomás-Barberán FA, González-Sarriás A, García-Villalba R, et al. (2017). Urolithins, the rescue of “old” metabolites to understand a “new” concept: Metabotypes as a nexus among phenolic metabolism, microbiota dysbiosis, and host health status. *Molecular Nutrition & Food Research* 61(1). doi:10.1002/mnfr.201500901. [REF:tomas-barberan2017]

Turroni S, Fiori J, Rampelli S, et al. (2016). Fecal metabolome of the Hadza hunter-gatherers: a host-microbiome integrative view. *Scientific Reports* 6:32826. doi:10.1038/srep32826. [REF:turroni2016]

Walsh L, Hill C, Ross RP (2023). Impact of glyphosate (Roundup) on the composition and functionality of the gut microbiome. *Gut Microbes* 15(2):2263935. doi:10.1080/19490976.2023.2263935. [REF:walsh2023]

Wang Z, Klipfell E, Bennett BJ, et al. (2011). Gut flora metabolism of phosphatidylcholine promotes cardiovascular disease. *Nature* 472(7341):57-63. doi:10.1038/nature09922. [REF:wang2011]

Wastyk HC, Fragiadakis GK, Perelman D, et al. (2021). Gut-microbiota-targeted diets modulate human immune status. *Cell* 184(16):4137-4153.e14. doi:10.1016/j.cell.2021.06.019. [REF:wastyk2021]

Wu GD, Chen J, Hoffmann C, et al. (2011). Linking long-term dietary patterns with gut microbial enterotypes. *Science* 334(6052):105-108. doi:10.1126/science.1208344. [REF:wu2011]

Zeb F, Wu X, Chen L, et al. (2020). Effect of time-restricted feeding on metabolic risk and circadian rhythm associated with gut microbiome in healthy males. *British Journal of Nutrition* 123(11):1216-1226. doi:10.1017/S0007114519003428. [REF:zeb2020]

Zeevi D, Korem T, Zmora N, et al. (2015). Personalized nutrition by prediction of glycemic responses. *Cell* 163(5):1079-1094. doi:10.1016/j.cell.2015.11.001. [REF:zeevi2015]

Zinöcker MK, Lindseth IA (2018). The Western diet-microbiome-host interaction and its role in metabolic disease. *Nutrients* 10(3):365. doi:10.3390/nu10030365. [REF:zinocker2018]

End of Chapter 20 draft.

Chapter 21: The Other Levers — How the Rest of How We Live Shapes the Ecosystem Inside Us

21.1 The Garden Has Weather

At the end of the previous chapter we left the reader with a practical promise: that the single most powerful daily lever most of us have over our microbiome is what goes onto the plate. Everything in Chapter 20 — the thirty-plant rule, the short-chain fatty acids, the mucus layer, the rediscovered importance of variety over quantity — was built on that lever. If the only message a reader remembers from this book is *eat more kinds of plants*, the book will have done its job.

But a gardener who tends the soil but ignores the weather will have a poor garden. And our microbiome, like any ecosystem, has weather. Some of it we control; some of it we do not. Some of it we invented in the last hundred years without realising we had invented a climate. This chapter is about those other inputs — the temperature of the room, the light on your eyelids at midnight, the dust that used to be on the furniture, the dog that isn't in the house, the stranger on the tram, the alarm clock that goes off when the rest of your organs would rather be asleep. Most of them are not things you think about at all. That is precisely why they are worth a chapter.

I should warn you at the outset that the evidence here is uneven. For some of the levers — artificial light and the circadian rhythm, for example, or the exposure of infants to household pets — the data are now strong enough to support firm statements. For others — what a HEPA filter in your bedroom might be doing to the airborne microbes you would otherwise inhale, or whether decades of living in a thermally constant twenty-two degrees has narrowed the range of microbial conditions your body experiences — we are still at the stage of reasonable mechanistic speculation rather than settled science. I will flag the difference as we go. Where I am guessing, I will say so. Where the literature is thin, I will say that too.

One more thing, before we start. Some of the topics in this chapter arrived in my notebook only very late in the writing of this book — not because they are unimportant, but because they fell outside the usual frame of microbiome research, which has been dominated by the diet–gut–immunity triangle. It was only after finishing Chapter 20 that I began to ask whether the ordinary domestic environment of a modern person — the air-conditioning that runs all day, the ceiling lights that burn until midnight, the HEPA filter that became standard after the pandemic — might itself be a quiet, continuous intervention on an ecosystem that evolved under conditions none of those appliances existed. I do not have all the answers. What I have, and what I will lay out in what follows, is a mechanistic case for why these things should matter, together with the slender but real body of published evidence that has begun to test the case directly.

This chapter is organised outward, roughly, from the body to the household to the outside world. We begin with the light that reaches your retina, move through the temperature of your bedroom and the air that circulates through it, step out into the dust and the pets and the people you live with, and end with the exercise, sleep, stress, water, and medications that complete the picture. Each section is a lever. None of them is as powerful as diet, considered alone. Considered together, they probably rival it.

21.2 The Forgotten Clock: Artificial Light and the Diurnal Microbiome

What the constant-darkness mice told us

Let me begin with an experiment that sounds trivial and turns out not to be.

In 2018, a team of researchers at Southern Medical University in Guangzhou took laboratory mice — the genetically uniform, biologically boring animals on which much of modern gut microbiome research is done — and divided them into two groups. One group lived in a normal twelve-hour light / twelve-hour dark cycle, the familiar laboratory default. The other group lived in constant darkness for the duration of the experiment. The mice were otherwise handled identically. They were fed the same food, housed at the same temperature, handled by the same technicians. The only variable was whether or not they could see the light come on in the morning [REF:wu2018].

The team then sampled the gut microbiota of both groups at multiple points across the day. In the normally lit animals, the composition of the gut microbial community rose and fell through the twenty-four-hour cycle. Different taxa dominated at different times of day. This diurnal oscillation of the gut microbiome is a real phenomenon, first described in detail by Christoph Thaiss and Eran Elinav in work we already met briefly in section 20.10 of the previous chapter [REF:thaiss2014] [REF:thaiss2016]. It is not an artefact of sampling; it is a genuine cycling of who-is-abundant-when inside the bowel.

In the mice kept in constant darkness, the oscillation collapsed. Almost every part of the intestine lost its rhythmic pattern. A bacterium called *Clostridia*, which in normally lit animals had a low, controlled presence in the small intestine, bloomed to several times its baseline abundance. The community was still there; it had simply forgotten what time it was. And when the community forgot what time it was, it lost the internal choreography that had kept certain species in check.

That last point is worth sitting with. We are used to thinking of the gut ecosystem as a collection of inhabitants, like animals in a jungle. The Guangzhou experiment adds a dimension most ecology textbooks do not bother with: the jungle has a clock, and when you break the clock, the tigers get loose.

The eye, the brain, and the bowel

If the experiment had only shown that constant darkness was bad for mice, we might reasonably have dismissed it as a laboratory curiosity. What makes the story interesting for humans is the circuit by which light, landing on the retina, actually reaches the gut.

In 2022, a group at National Taiwan University published a paper in *EMBO Reports* that asked, in effect: how does the light on the eye talk to the bacteria in the bowel? They already knew that circadian misalignment disturbed the gut microbiota, but the intermediate steps had not been worked out. Using genetically engineered mice that lacked specific light-detecting cells in the retina, they showed that a small, specialised population of retinal cells called *intrinsically photosensitive retinal ganglion cells* — ipRGCs for short — was responsible for transmitting day-length information from the outside world to the circadian pacemaker in the brain, and from there to the intestinal community [REF:lee2022].

The ipRGCs are an odd and beautiful class of cell. They are neither rods nor cones; they do not contribute to the image you see. They carry a photosensitive pigment called melanopsin, and their

job is to tell the suprachiasmatic nucleus — the brain’s master clock — what time of day it is. They are the biological equivalent of the watch face on a diver’s wrist that glows under water. You do not use them to read. You use them to know.

The Taiwanese group showed that when the ipRGCs were silenced, the diurnal oscillation of the gut microbiota broke down in the same way Wu and colleagues had seen with constant darkness. They then did something else. They left the ipRGCs intact but exposed the mice to dim light at night — the kind of low-level domestic light spill that comes from a television set in another room, or a street lamp shining through thin curtains, or a hallway nightlight — and showed that this was enough to alter both the composition and the normal daily rhythm of the gut community. Not bright light. Not flashbulbs. The kind of light most of us do not bother to switch off because we have never considered that it might be doing anything.

Independently, in the same year, John Brooks and colleagues in Lora Hooper’s laboratory at UT Southwestern published a paper in *Cell* that completed a different arm of the same picture. They showed that the microbiota itself generates diurnal rhythms in innate immunity, through segmented filamentous bacteria that attach and detach from the epithelium according to the host’s feeding schedule, which is in turn entrained by the light cycle. The host’s antimicrobial protein production — the first line of defence against any swallowed pathogen — rises and falls with the feeding rhythm and therefore, indirectly, with the light [REF:brooks2021]. So light, via the retina and the brain, regulates the rhythm of eating; eating regulates the rhythm of bacterial attachment; bacterial attachment regulates the rhythm of antimicrobial defence. Pull one thread and the whole quilt follows.

What this means for someone who lives in a lit building

Most of us do not live under thirty-six-hour fluorescent tubes. Why does any of this matter for the reader of this book?

It matters because, for the first time in several hundred thousand years of human existence, a substantial fraction of the population is now under artificial light for a large, continuous block of what used to be night. A streetlight through a bedroom curtain. A blue-screened phone held thirty centimetres from the face in the dark. A bathroom fluorescent that is on every time a child wakes up to go to the toilet. A kitchen under-cabinet LED that is never quite off. Modern domestic light is low in intensity by outdoor standards, but it is several orders of magnitude brighter than the upper limit under which ipRGCs can reliably decide it is night — which, experimentally, is around five lux, roughly the level of a full moon seen from a clearing.

Shift workers are the obvious human test case — the clearest example of chronically mis-timed light exposure. In a 2021 study published in *PeerJ*, Ann Rogers and colleagues at Emory University recruited fifty-one nurses working twelve-hour day or night shifts, collected stool samples from them before and after a rotation, and sequenced the gut microbiota of each [REF:rogers2021]. The differences they found between day-shift and night-shift nurses were, perhaps surprisingly, modest — small shifts in beta diversity, a handful of operational taxonomic units that differed between groups. But the signal was there, and the more recent review literature collating shift-work and gut-microbiota studies [REF:lopezsantamarina2023] consistently points in the same direction: shift workers show altered gut microbial communities that, on current evidence, are probably one of the mechanisms by which rotating and night shifts raise the risk of metabolic disease. Shift workers are the extreme; but most of us are, in a quieter way, living a very mild version of the same experiment. We eat at times our grandparents would have found shocking. We use screens in bed. We put

babies in rooms that are technically dark but lit enough for a parent to find their way to the cot. In aggregate, we have created, in the space of about a century, an indoor lighting environment that no previous generation had, and that was not present during the period when our gut microbes were assembling their rhythms.

Hong and colleagues, in 2020, took this one step further in a mouse experiment at Beijing Normal University. They exposed mice to constant light — the modern city in caricature — and watched them become obese and insulin-resistant alongside a measurable shift in their gut microbial composition. Giving the mice supplementary melatonin (the hormone the suprachiasmatic nucleus releases at night, and which modern light spill suppresses) partly rescued the metabolic disturbance *and* rescued the gut dysbiosis [REF:hong2020]. A constant-light environment, a night-lowered melatonin signal, a disturbed microbiota, a metabolic penalty: in the mouse these now connect in a single experimental chain.

How strong is the human evidence?

Less strong than the animal evidence, and I want to say that clearly. We do not yet have a large randomised trial that switches human bedroom lighting from “dim at night” to “proper darkness” and measures what happens to the gut microbiome over a year. What we do have are shift-work observational studies, the mechanistic chain described above, and a small but growing body of human work on circadian misalignment and gut bacterial rhythmicity. The most honest summary is: the direction is clear, the magnitude in a domestic setting is not yet known, and the practical advice that falls out of this is sensible regardless of how microbial the final answer turns out to be.

That practical advice is approximately what your grandmother would have told you. The bedroom should be dark — properly dark, the kind of dark where you cannot see your hand until your eyes adjust. Heavy curtains or an eye mask are cheap interventions. Hallway nightlights should use red or amber LEDs, because the ipRGCs are relatively insensitive to long wavelengths; a red light at two in the morning does not tell the master clock that dawn has come. Phones should not live on the pillow. And — this is the one most people will find hardest — the room you eat in, and the time at which you eat, should be synchronised with the light your retina is seeing. If you eat your largest meal at nine in the evening under bright kitchen lights, you are presenting the brain with a mixed signal — bright light says morning, the stomach says evening — and the internal clocks begin to disagree about what time it is. That disagreement, as we saw in Chapter 20’s discussion of circadian misalignment, is itself bad for the gut community [REF:thaiss2014].

In a book about microbes, it is a strange thing to end a section with advice about curtains. But the lever is real, it is almost free, and — unlike most things in this chapter — it does not require you to add a thirty-first plant to your shopping list.

21.3 Twenty-Two Degrees All Year: The Thermostat and the Unfelt Season

What our bodies used to do that they no longer do

The most reliable feature of the climate inside the average modern home is that it has no seasons. In winter the room is warm. In summer the room is cool. The variance is narrow — somewhere around twenty to twenty-four degrees Celsius in almost any month of almost any year, in almost any dwelling in the developed world. This is historically extraordinary. Until central heating and

air-conditioning became standard, roughly in the middle of the twentieth century, the human body experienced a temperature cycle across the year that was at least fifteen or twenty degrees wide — and often much more, if you lived in a climate with real winters — and a daily temperature cycle that was considerably wider than today's, because houses cooled overnight and warmed at dawn as the sun rose and the stove was lit.

This matters to microbes for two reasons, one of which is reasonably well established and one of which is speculative.

The reasonably well established reason is that the mammalian body, when exposed to cold, activates a tissue called *brown adipose tissue*, or brown fat, which burns calories to generate heat. Brown fat is a metabolic organ in its own right: it can be dialled up by the sympathetic nervous system in response to cold, and when it is dialled up, it changes what the whole body does with food. The research here is mostly in mice, but the mouse data are very clear. When mice are exposed to environmental cold for extended periods, their gut microbiota change composition substantially. Earlier papers had argued that the gut microbes are actually *part* of the thermogenic response, not just a bystander to it; a 2020 paper in *Cell Metabolism* by Krisko and colleagues at Weill Cornell pushed back carefully, showing that cold adaptation and energy expenditure in germ-free and microbiome-depleted mice are essentially intact, and concluding that the microbiome is *dispensable* for adaptive thermogenesis itself — though they still found that the microbial community does change with cold and does contribute to other metabolic outputs (specifically, to amino-acid metabolites that feed hepatic gluconeogenesis) [REF:krisko2020]. The literature is arguing, but what is not disputed is that ambient temperature changes the gut community.

The speculative but mechanistically plausible reason is subtler. Many of the microbes that live on our skin, in our airways, and — to a lesser but real extent — in our gut have temperature preferences. Temperature is not a minor variable for a bacterium; it is one of the core determinants of whether enzymes fold correctly and whether cell membranes stay fluid. A *Staphylococcus epidermidis* cell on your forearm is optimally active at around thirty-two to thirty-four degrees, which is why skin, rather than the deep body, is where it lives. If the skin surface temperature changes (because the room has changed), the balance between a slightly warmer skin taxon and a slightly cooler one may shift. The same argument applies with different absolute numbers for the nose, the throat, and even the mucosal surfaces of the lower bowel, although the core body temperature buffers the bowel against most of this.

The speculative part is that no one has yet done the study that would tell us what a lifetime of thermostat-stable twenty-two-degree indoor living has done to the human skin and airway microbiome, compared with a lifetime of ten-degree winters and thirty-degree summers in an unheated house. I suspect it has done something; I cannot cite a paper that proves it yet. I raise the issue here because the absence of evidence is itself worth knowing about, and because if you are going to design such a study, you need to have noticed the question first.

What I can say with some confidence is that *feeling cold occasionally* — that is, not treating climate control as a moral imperative — is unlikely to do harm and may do good. The thermogenic and metabolic effects of mild cold exposure are well documented in human physiology, even though the microbiome piece of that story is still being written. A slightly cooler bedroom at night, in particular, has additional benefits for sleep quality that are not speculative at all, and sleep quality, as we are about to see in section 21.9, is itself a microbiome variable.

Humidity: the unspoken second axis

The other thing air-conditioning does that is easy to forget is dry the air. A cooling unit pulls moisture out of the air as a necessary side-effect of chilling it, and a heating unit dries the air because warmed air holds more water vapour at the same absolute humidity. The winter bedroom in a heated house is typically at around twenty to thirty percent relative humidity — the equivalent of a desert. The summer bedroom in an air-conditioned house is often at around thirty to forty percent — still dry by outdoor standards in most temperate climates.

The skin microbiome, the mycobiome of the airways (*Malassezia*, *Candida*, *Aspergillus* species we met in Chapter 10), and the overall fungal load of the indoor environment all respond to humidity. Low humidity dries mucus membranes, which in turn impairs mucociliary clearance — the escalator that carries inhaled microbes out of the airways. Seasonal variation in humidity is one of the reasons respiratory infections spike in winter in heated buildings and again in air-conditioned offices in summer. The observation is decades old; what is newer is the growing recognition that the ecology of the built environment — the set of microbes that actually live in and on the fabric of our homes — is shaped by our heating and cooling systems to a degree we have not fully reckoned with.

For the practical reader, a humidifier set to around forty to fifty percent relative humidity in a heated winter bedroom, and a humidifier or damp cloth in an air-conditioned summer one, is a cheap intervention with a good safety record. The microbiome benefit is inferred from mechanism rather than proven. The benefit to mucous membranes is not in dispute at all.

21.4 The Invisible Garden in the Air: HEPA Filters, UVC Lamps, and the Post-Pandemic Reflex

Cleaner than what?

When SARS-CoV-2 ripped through the world in 2020 and 2021, most of us — doctors included — learned a new word: *aerosol*. We learned that a cough, or a breath, or even a long conversation, produces a plume of tiny droplets that can stay suspended in still indoor air for minutes and can, under the right conditions, carry a virus. The response, over the months that followed, included a surge in sales of portable air purifiers — typically units containing a high-efficiency particulate air filter, known as a HEPA filter, which can capture particles down to about three tenths of a micrometre with impressive efficiency. In many workplaces and a surprising number of homes, these units now run continuously. Some models add an ultraviolet-C lamp to kill captured microbes. Some add ionisers. Some advertise “germicidal action” with all three.

This is a book about microbiology, so I want to be careful here. HEPA filtration, properly installed, is an effective and valuable public-health intervention in settings with a real airborne infectious disease risk — hospital isolation rooms, tuberculosis clinics, airline cabins, and laboratories working with dangerous pathogens. I am not going to argue against using them where they are indicated. What I do want to examine is a different question, one that almost nobody thought to ask during the pandemic: if a HEPA filter removes the viruses and bacteria you do not want from the air, does it also remove the microbes you might have been healthier for inhaling?

It is not a silly question. Several lines of evidence converge on the idea that airborne microbial diversity, particularly the environmental bacteria and fungi derived from soil, plants, and untreated outside air, contributes to the maturation and tolerance of the immune system, especially in early

childhood — an argument we met briefly in the hygiene-hypothesis discussion of Chapter 6 and will return to later in this chapter. The question of what continuous HEPA filtration does to those environmental microbes in a living room has, as far as I can find, essentially never been studied as a microbiome outcome in humans.

Let me be honest about what the evidence actually shows, because this section would be easy to get wrong. What the literature does *not* yet contain is a clean randomised trial in which the continuous running of a domestic HEPA unit in a home bedroom is compared with no filtration, and the airborne and gut microbiomes of the occupants are tracked over months or years. I searched for this paper while drafting this chapter and could not find it. The reader should know that the specific question “does a HEPA unit in the bedroom of a healthy child for a year change their gut or airway microbiome?” has not, to my knowledge, been directly answered.

What the literature *does* contain is an accumulation of indirect evidence. Studies of the built-environment microbiome have repeatedly shown that the airborne microbial communities of indoor spaces are a mixture of human-derived material (skin flakes, the aerosols of speech, cough droplets) and outside-derived material (pollen, soil bacteria, the airborne fraction of the external weather). Studies sampling the same rooms before and after the opening of a window, or before and after the introduction of a dog, have consistently shown that both of these interventions *increase* the diversity of airborne bacteria in the indoor space and move the community towards the composition of the outside air. Continuous HEPA filtration, in parallel, removes the particulates that carry outside-derived microbes into the room as well as those that carry pathogens, and on first principles there is no reason to believe that a filter capable of removing virus-laden droplets is not also removing environmental bacteria attached to particles of the same size range.

What the filter actually collects

There is a separate and almost poetic body of work in which researchers have taken the HEPA and HVAC filters *out* of buildings and read the DNA off what the filters trapped. This is called *filter forensics*, and it has become a quiet cottage industry in environmental microbiology. A 2018 paper in *Microbiome* by Maestre and colleagues, for instance, showed that residential HVAC filters captured rich, recoverable bacterial and fungal DNA reflecting the occupants and the outdoor environment of the home [REF:maestre2018]. A separate line of work by Miles Richardson, Simon Lax, and Jack Gilbert in dormitories showed that the microbial signatures left behind by individual cohabitants were distinctive enough to identify *who* had been in a shared space, even after repeated mixing [REF:richardson2019]. The filter is not just catching particles; it is archiving the microbial history of the people, pets, and outdoor air that passed through it. Some of what it captures is nuisance taxa and pathogens. Some is probably part of the normal traffic of the human ecosystem — the shared skin bacteria of cohabiting couples, the airborne fraction of the oral microbiome that lands back on surfaces and gets re-inhaled, the commensal microbial cloud that each person carries wherever they go. An ultra-efficient filter running twenty-four hours a day does not discriminate between the harmful and the benign. It just removes.

A tentative verdict

Here is where I want to be honest about the limits of what we know.

I am *not* saying: throw out your air purifier. In settings where airborne infectious disease risk is real — an immunocompromised family member, a viral epidemic, a workplace with known exposures

— filtration is valuable and I will recommend it to my own patients in those situations. I am not saying anything that contradicts the known value of HEPA filtration in those settings.

What I *am* saying is that the post-pandemic reflex of running a HEPA unit continuously in the bedroom of a healthy adult or child, as if “cleaner air” were unambiguously better, is not a policy the microbiome literature directly supports — and, more importantly, not a policy the microbiome literature has properly examined. The airborne microbiome of an open-window home with a dog is almost certainly richer and more diverse than the filtered air of the same room with the window closed and the filter running, and — on the weight of the hygiene-hypothesis evidence we will revisit in section 21.5 — the diverse air is more likely to be the one that supports healthy immune development. But the reader should hear the qualification as well as the claim. The direct experiment has not been done. If you are healthy, your dog is healthy, and no one on your street has tuberculosis, the most evidence-backed air intervention you can make in your own home is to *open the window for fifteen minutes a day*, regardless of the season, and to let the weather inside — and that recommendation rests less on the HEPA-versus-no-HEPA literature, which is thin, than on the considerable literature we will discuss in section 21.5 showing that contact with outdoor and environmental microbial diversity is protective against allergic and autoimmune disease.

UVC lamps are a separate and more worrying category. Some of the air purifiers now marketed to consumers include a UVC stage claimed to inactivate bacteria and viruses passing through the unit. UVC light in the 254 nanometre band is indeed germicidal; the question is whether the dose delivered to air passing through a domestic unit at typical flow rates is sufficient to inactivate an appreciable fraction of the organisms, and the answer varies enormously between units and is almost never verified independently. More concerning, UVC at certain wavelengths can generate ozone inside the unit, and ozone at the concentrations produced by some poorly designed home units is an irritant to the lung mucosa and has its own effects on the airway microbiome that are unlikely to be positive. For the domestic consumer I would not currently recommend UVC air units as a first-line intervention; the risk-benefit ratio is not favourable for healthy people, and the microbiome literature does not yet provide a reason to change that assessment.

Ionisers are similar. Bipolar ionisation was heavily marketed during the pandemic as a way to inactivate airborne viruses; independent testing has been mixed to unfavourable, and the by-products of the process include small amounts of ozone and various reactive oxygen species. I would not run one continuously in a bedroom.

If the reader takes one thing from this section, it should be this: in a healthy home, *air variety* is probably more important than *air purity*, and the cheapest way to achieve variety is to let the outside in. Opening the window, for a quarter of an hour, at least once a day, in almost any weather, delivers more microbial diversity to the home than almost any other intervention. It is also the intervention most at odds with how the modern house is designed and run, which is why it is worth making explicit.

21.5 Dirt, Biodiversity, and the Old Friends We Stopped Meeting

If there is one finding in environmental microbiome research that has survived twenty years of replication attempts, it is this: children who grow up with more microbial exposure early in life have less allergic disease, less asthma, and less autoimmune illness later in life. The observation is old; the mechanism is newer; the policy implications are still being argued over. But the core of the

story has become difficult to dismiss.

The foundational paper for most modern thinking on this question was published in the *New England Journal of Medicine* in 2016 by Michelle Stein and colleagues working in the American Midwest [REF:stein2016]. Their study compared two religious farming communities — the Amish and the Hutterites — who are genetically similar (both descended from a small founder population of German-speaking Anabaptist farmers), share similar diets low in processed foods, have large families, and practise lifestyles at the conservative end of the modern spectrum. On almost every variable that epidemiologists track, the two groups look alike. The Amish and the Hutterites should therefore have similar rates of childhood asthma.

They do not. Asthma among Amish schoolchildren runs at about 5%; among Hutterite children it is around 21%. The gap is enormous, consistent across years, and has been confirmed in multiple independent samples. Given how similar the populations are, something about the way the two communities farm must be responsible, and that something had to be exposure-related rather than genetic, because the two groups share much of their ancestry.

The difference, as Stein and colleagues showed, is that Amish families farm in the old way — small, mixed farms where the barn sits a few metres from the kitchen door, where children run in and out of the cowshed from the time they can walk, where hay is pitched by hand and milk is hand-drawn, and where household dust accumulates a thick and constant inoculum of barn-derived microbial material. Hutterite families, by contrast, have modernised: they farm at industrial scale, the barns are separate from the living quarters, children are kept out of the animal areas, and the homes are cleaner. When the researchers collected household dust from both groups and measured its endotoxin content (endotoxin is a marker of Gram-negative bacterial debris and a reasonable proxy for overall bacterial load), Amish house dust contained roughly seven times more endotoxin than Hutterite dust.

What happens when you take that Amish dust and blow it up the noses of laboratory mice? Stein's group did exactly this experiment. Mice exposed repeatedly to aerosolised Amish dust developed robust innate immune activation in their airways, with high numbers of neutrophils producing the right cytokine pattern; when these mice were then challenged with a standard asthma-provoking allergen, they were strikingly protected from the airway inflammation their unexposed littermates developed. Hutterite dust, given in the same way, did not confer the same protection. Something about the composition or quantity of microbial material in Amish dust was reprogramming the innate immune response of the airway in a way that suppressed allergic inflammation.

The Amish-Hutterite study is not the only piece of evidence for this relationship; it is simply the most elegant natural experiment we have, because it holds so many potentially confounding variables constant. Similar observations come from comparisons between Finnish and Russian Karelians (genetically related populations separated by a border and by dramatically different levels of hygiene, with more allergy on the Finnish side), between children born on European farms and children born in European cities in the ALEX and PARSIFAL cohorts, and — most recently — between rural and urban populations in rapidly urbanising Asian countries. Across all of these, a consistent pattern emerges: exposure to a rich, stable, earthy microbial environment in early life is protective against allergic and autoimmune disease, and the effect is strongest when the exposure occurs in the first year or two of life.

The most developed theoretical framework for understanding why this should be so is the “Old Friends” hypothesis, proposed and developed over two decades by the British immunologist Graham Rook [REF:rook2013] [REF:rook2014]. Rook's argument is, at root, an evolutionary one. For

almost all of human history — and for the hundred million or so years of mammalian history before us — our immune systems have co-evolved in the presence of a particular set of environmental microbes and macro-parasites: the soil-dwelling saprophytes that colonise our food and water, the environmental mycobacteria that live in damp soil and decaying vegetation, the helminth worms that inhabited almost every human gut for most of our species' existence, and the commensal microbes passed around by the large extended families and the farm animals with which humans lived in continuous close contact. These are the “Old Friends” — not pathogens, not innocent bystanders, but *tonic* microbial and parasitic inputs that the immune system came to rely on for its calibration.

Rook's proposal is that these inputs did a specific immunological job: they continuously trained the regulatory arm of the immune system to tolerate innocuous antigens. Regulatory T cells (Tregs), the immune cells we met in chapter 12 as peacekeepers that suppress inappropriate inflammatory responses, require environmental signals to differentiate and expand properly. In a world rich in Old Friends, this regulatory compartment is constantly being exercised and replenished; in a world stripped of them — a world of paved streets, water-treated taps, chlorinated swimming pools, sealed houses, single-child families, and early antibiotic exposure — the regulatory compartment never develops its full repertoire. The immune system, deprived of its usual tuning signals, becomes prone to mistakes. It misidentifies harmless pollen as a threat (allergy), attacks its own tissues (autoimmunity), or produces chronic low-grade inflammation in response to nothing in particular.

The Old Friends framework is *not* the older and clumsier “hygiene hypothesis” that David Strachan proposed in 1989, which suggested that childhood infections protect against allergy and which was widely misread as an argument for deliberately exposing children to germs. Rook's revision is subtler and more biologically specific. The protective microbes are not childhood pathogens; they are the ancient, tolerated, environmental partners that modern life has stripped out of the human environment. The remedy is not to catch more colds; the remedy is to restore contact with soil, animals, and diverse environmental microbial communities — and to do so, as much as possible, in early life, when the immune system is still being assembled.

From a book-writing point of view, I want to be careful here, because the Old Friends hypothesis is sometimes stretched further than the evidence can carry it. Not every increase in allergy and autoimmunity in the industrialised world is explained by loss of microbial exposure; diet, pollution, genetics, improved diagnosis, and many other factors contribute. The hypothesis is best regarded as *one important piece of a larger puzzle*, not as a complete explanation. What makes it worth taking seriously, rather than dismissing as another appealing just-so story, is that it has generated concrete predictions that have been tested and confirmed — most notably in the Amish-Hutterite work — and that it is biologically coherent with everything we now know about how the microbiome talks to the immune system.

What does this mean for the reader sitting in a modern home? A few things. It means that the cumulative effect of the design choices we have discussed so far — sealed windows, thermostat-held climate, HEPA-filtered air, paved yards, pet-free apartments, chlorinated water — is not just a loss of microbial variety in the abstract; it is a systematic removal of the environmental signal that the immune system was built to expect. It means that the cheapest and most evidence-backed interventions to restore some of this signal are profoundly low-tech: a garden with real soil, a dog, an open window, time spent in forests and fields, a willingness to let young children get dirty. It means that the more we understand about the Old Friends relationship, the more we come to see the modern sanitised home not as the neutral zero-point from which departures must be justified, but as itself a *departure* — a recent and poorly tested experiment in which the human animal has been lifted out of the environment its immune system expects, and placed into a microbiologically

impoverished one. The question “is this exposure safe?” therefore needs to be supplemented by its mirror: “is the *absence* of this exposure safe?” Often, the honest answer to the second question is: we do not know, and the emerging evidence suggests that it is not.

21.6 The Dog in the Doorway: Household Pets and the Infant Microbiome

If you asked a microbiome researcher to name the single intervention a young family can make, in the first year of their child’s life, that is most likely to meaningfully broaden the child’s developing gut flora, the answer — increasingly, across independent cohorts — is: *get a dog*.

The cleanest evidence for this comes from the Canadian Healthy Infant Longitudinal Development (CHILD) cohort study, a prospective birth cohort that enrolled over three thousand Canadian infants and collected faecal samples, household environment data, and detailed health information at regular intervals from birth through early childhood. In a 2017 analysis, Hein Min Tun and colleagues used CHILD data to ask a specific question: does the presence of a furry pet in the household during pregnancy and early infancy alter the infant gut microbiome measurably, and in a direction that might plausibly matter for later immune outcomes? [REF:tun2017]

Their answer was unambiguously yes. Infants raised in households with at least one furry pet (most commonly a dog, but cats also contributed) had significantly higher relative abundance of two bacterial genera — *Ruminococcus* and *Oscillospira* — that in other studies have been linked to reduced risk of childhood allergic disease and reduced risk of obesity. The effect was dose-dependent, in the sense that exposure both during pregnancy and after birth had a larger effect than exposure at only one of those periods, and it persisted after adjustment for breastfeeding, delivery mode, antibiotic exposure, and the usual confounders. Most strikingly, the pet effect was strong enough to partly *compensate* for the microbiome impoverishment normally seen after caesarean delivery or after antibiotic exposure in the perinatal period: C-section babies raised with a dog in the home had gut microbiomes that looked more similar to vaginally-delivered babies raised without a dog than to their own C-section peers in pet-free households.

Independent cohorts have broadly confirmed the direction of the effect, even if the specific taxa that shift vary between studies. A 2020 paper by Nielsen and colleagues [REF:nielsen2020] used the Danish COPSAC2010 cohort to show that pet exposure in early infancy was associated with increased infant gut microbiome diversity and with altered levels of specific short-chain fatty acid producers, and that these changes mediated some of the already-known protective association between pet exposure and later allergic sensitisation. A German cohort, a Dutch cohort, and multiple smaller studies have produced broadly similar findings. The detail differs; the direction of the effect does not.

Mechanistically, several things are going on, and they overlap. Dogs and cats drag the outside world into the house on their paws and fur — they carry soil, grass, pollen, and the microbial communities of wherever they have been — and they transfer this material to floors, furniture, and the hands and faces of the children who crawl and toddle over that terrain. Dogs in particular share a portion of their own skin and oral microbiome with the humans they live with; studies of household microbial sharing (which we will get to in the next section) have found that dogs behave, microbiologically, like an extra cohabitant whose effect on household flora is larger than any single human’s. And finally, dogs structure the behaviour of the family: owning a dog means walking outside in weather the family might otherwise avoid, which in turn means incidental contact with the full seasonal

range of environmental microbes.

The practical implication is interesting, and — importantly — it is one of the very few microbiome interventions that is essentially free of downside for a family that was already disposed to considering a pet. The Old Friends hypothesis predicts that dogs should be beneficial; the CHILD cohort and its siblings have confirmed that they are; the effect is robust across populations; and there is no plausible mechanism by which routine exposure to a healthy pet *harms* the infant microbiome. Allergic sensitisation to the pet itself is, counter-intuitively, less common in children exposed from birth than in those who meet a dog for the first time at age five. If a family is considering a pet and a young child is in the picture, the microbiome literature now points fairly clearly in the direction of saying yes, bringing the pet in early, and letting the child and the dog share a considerable amount of physical contact.

A word of caution is in order, because nothing about this recommendation should be read as license to expose vulnerable people to genuinely sick animals, or to overlook zoonotic risks. Immunocompromised patients, pregnant women handling cat litter (because of toxoplasmosis), and anyone receiving intensive chemotherapy face real trade-offs that are more complicated than the average healthy family does. For the typical healthy household, however, the risk-benefit calculus favours pet ownership rather clearly, and the microbiome data have now added a specific reason to the older, more general reasons — companionship, exercise, childhood wellbeing — that people keep dogs.

21.7 The Household as a Microbial Unit: What We Share With the People We Live With

When two strangers move into a flat together, they bring two different microbiomes with them. Within a few weeks, if they share meals, bathrooms, couches, and doorknobs, a curious thing starts to happen: their microbiomes begin to converge. Not completely, and not on every body site equally, but enough that investigators sampling the skin, nose, and gut of cohabitants can distinguish “same household” from “different household” with surprising accuracy on the basis of microbial sequences alone.

The most comprehensive early documentation of this phenomenon came from the “Home Microbiome Project” led by Jack Gilbert at Argonne National Laboratory, published in *Science* in 2014 by Simon Lax and colleagues [REF:lax2014]. The study followed seven families (eighteen people in total, plus three dogs and a cat) over six weeks, swabbing skin surfaces, kitchen and bathroom fixtures, doorknobs, and floors every few days. The results were striking. Each household developed a characteristic microbial “signature” — a pattern of species and their relative abundances — that was far more similar to itself over time than to any other household. Within a household, there was constant microbial exchange: the floor microbiome and the shoe microbiome tracked the outside environment; the countertop microbiome tracked the skin and hand microbiomes of the residents; the doorknob and light-switch microbiomes were essentially human-skin-derived and distinguishable by person. When a family moved house during the study, the new house acquired the old house’s microbiome within a day — the families brought their microbial community with them, and it overwrote whatever had been there before.

A separate line of work — most influentially from Brent Stephens and colleagues at Illinois Institute of Technology, and from the Gilbert lab — has shown that this household microbiome is not simply

the sum of the people who live in it. It has its own ecology. Dust in the corners of rooms accumulates over weeks a stable community that includes skin shed from residents, environmental bacteria tracked in on shoes, outdoor air input via the windows and HVAC, and, in humid bathrooms, locally-blooming moisture-loving genera. Vacuuming, mopping, and airing the house reshape but do not eliminate this community; within a few days of a thorough cleaning, the dust flora returns to something close to its previous composition.

More recently, attention has turned from the *built* microbiome to direct person-to-person sharing of gut microbes within households. In a major 2023 *Nature* paper, Mireia Valles-Colomer and colleagues from the Segata lab in Trento analysed nearly ten thousand metagenomes from cohabitants in multiple countries to reconstruct strain-level transmission of gut bacteria between people [REF:vallescolomer2023]. Their results show that vertical transmission from mother to infant is the dominant source of early-life gut strains (as expected), but that horizontal transmission between cohabitants continues throughout life at levels much higher than previously appreciated. Spouses who live together for years share a substantial proportion of their gut bacterial strains — far more than random pairs from the same population — and the degree of sharing increases with duration of cohabitation. Children share with their parents; siblings share with each other; in multi-generational households, grandparents share with grandchildren.

This has several important implications, and I want to spend a moment on the two that I think matter most for the general reader.

The first is that the gut microbiome is not a private possession, sealed within an individual body. It is a shared domestic ecology, maintained collectively by the people who live under the same roof. Your spouse's microbiome is, in a real sense, partly yours. The father who makes a radical dietary change in mid-life is not only remodelling his own gut community; he is, slowly, exchanging strains with everyone else in the household. The child who is put on a course of antibiotics does not suffer the disturbance in isolation; the household reservoir of bacterial strains is part of what the child's gut will draw on as it recovers. This interconnectedness gives a new meaning to the old public health observation that health outcomes cluster within families. It turns out that one of the mechanisms by which lifestyle diseases and health behaviours run in families is, almost certainly, microbial.

The second implication is more uncomfortable. The widespread assumption that microbiome-focused dietary interventions — more fibre, more fermented foods, the changes we discussed at length in chapter 20 — are an individual choice with individual consequences is not quite right. The household is the unit; the rest of the household both helps and hinders the individual's efforts. A partner who refuses to eat vegetables and lives almost entirely on ultra-processed food is, in small but measurable ways, reintroducing the strains associated with that way of eating back into the household reservoir, from which they can re-colonise the person who is trying to eat better. Household-level interventions — where a whole family shifts together — would be predicted by this biology to be more effective than individual-level interventions. That is also what the (admittedly limited) interventional evidence tends to show.

The household microbiome framework does not lead to any particularly drastic recommendations. It does, however, re-frame a set of questions that we usually treat as individual questions into collective ones. It suggests that the microbiomes of the people we live with are part of our own microbial environment, and that caring for our own gut flora is, in practice, partly a matter of caring for the domestic ecosystem we share with them.

21.8 Moving the Body: Exercise as a Microbial Intervention

Exercise affects almost every physiological system we know how to measure, and in the past decade the microbiome has been added to that list. The question is no longer whether physical activity changes the gut microbiome — it does — but whether the changes are clinically meaningful, whether they explain any of exercise’s well-known benefits, and whether the effect is independent of the dietary and lifestyle shifts that usually accompany a decision to start exercising.

The early, eye-catching evidence came from a 2014 study of elite Irish international rugby players by Siobhan Clarke and colleagues at University College Cork [REF:clarke2014]. The researchers sequenced the gut microbiomes of forty professional rugby players during a pre-season training camp and compared them with matched sedentary controls. The athletes had markedly more diverse gut microbiomes, with higher relative abundance of *Akkermansia muciniphila* (a mucus-dwelling bacterium associated with metabolic health), more short-chain fatty acid producers, and a generally “healthier-looking” community profile. The problem with this study, which the authors openly acknowledged, is that elite rugby players are not just more active than the controls; they also ate dramatically more food, consumed more protein, had higher creatine intakes, and differed in dozens of ways that could have driven the microbiome differences independent of the exercise itself. The headline finding was suggestive but confounded.

The interventional evidence needed to address that confound came later. The cleanest trial I am aware of is a 2020 randomised controlled trial by Mika Motiani and colleagues in Finland [REF:motiani2020], which enrolled sedentary adults with insulin resistance — a group likely to benefit from any genuine metabolic effect of exercise — and randomised them to either sprint interval training or moderate-intensity continuous training for two weeks, with dietary intake carefully held constant. Both exercise protocols produced measurable changes in the gut microbiome: decreased *Bacteroidetes* relative to *Firmicutes*, increased *Akkermansia*, and increased plasma short-chain fatty acid levels. The changes correlated with improvements in whole-body insulin sensitivity, and importantly, they occurred without any change in diet. This was, as best we can tell, the direct effect of the exercise itself on the gut ecosystem.

The mechanisms by which exercise moves the microbiome are still being worked out, but several plausible routes have been identified. Exercise raises core body temperature modestly, alters gut blood flow and intestinal transit time, increases bile acid secretion, reduces systemic inflammation, and triggers the release of exercise-responsive signalling molecules (myokines from muscle, and, in longer sessions, stress hormones) that reach the gut lumen and may reshape which bacteria thrive there. In parallel, the act of breathing deeply and being outdoors while exercising brings the airway microbiome into contact with a wider range of environmental microbes than sitting indoors does — which means that, at least for outdoor exercise, some of the microbiome benefit may come through the same “Old Friends” pathway as gardening or dog-walking does.

What dose of exercise moves the needle? The current literature does not allow a precise prescription, but the pattern across studies is that it does not take much. The Motiani trial produced measurable effects in two weeks at moderate intensity. Observational studies comparing physically active adults with sedentary ones find microbiome differences at activity levels well below what an elite athlete does — levels closer to the standard public health recommendation of roughly 150 minutes of moderate exercise a week, or about twenty to thirty minutes most days. The effect is larger with outdoor exercise than with treadmill exercise (for the environmental exposure reason), and it appears to plateau rather than increase linearly: there is no evidence that doing twice as much produces twice the microbiome benefit.

For the purposes of this chapter, exercise is interesting less as a surprise and more as a reinforcement of the central theme: the levers that most reliably move the gut microbiome in a favourable direction are not exotic. They are the same levers that human bodies have always needed for health — movement, contact with the outside world, varied food, sleep, social and family connection. The microbiome literature is not overturning those familiar recommendations; it is providing a new and specific biological reason for why they matter.

21.9 Sleep and Stress: The Nervous System's Dialogue With the Gut

Few subjects in popular microbiome writing are more vulnerable to overclaim than the effects of sleep and stress on gut flora. The relationship is real; the mechanisms are being worked out; and at the same time, some of the more enthusiastic claims in the self-help literature rest on animal data that have not yet been replicated in humans, or on small observational studies whose direction of causation is impossible to establish. I will try to describe what the evidence actually supports.

Sleep first. The clearest experimental evidence that disrupted sleep reshapes the gut microbiome comes from shift-work studies and from short-term sleep-deprivation protocols. In 2016, Christian Benedict and colleagues in Sweden ran a small crossover trial in which healthy young men spent two nights either sleeping normally or being kept awake through the night, with identical diets provided in both conditions [REF:benedict2016]. Two nights of partial sleep deprivation produced measurable shifts in the ratio of *Firmicutes* to *Bacteroidetes* that, in the broader literature, have been associated with metabolic dysfunction. The changes were small, the study was short, and whether the shift would persist or reverse with habitual poor sleep is not known from that design alone. But the signal was real and has been replicated in larger observational studies of chronic short sleepers and shift workers.

A more recent and larger piece of work is a 2023 analysis by Sun and colleagues [REF:sun2023sleep] using data from a population-based cohort that paired detailed sleep metrics (from validated questionnaires and, in a subset, from actigraphy) with gut metagenomic data. Habitually poor sleepers and those with fragmented sleep had measurably different gut community structures from good sleepers, with reduced alpha diversity, lower relative abundance of several butyrate-producing genera, and higher relative abundance of taxa previously linked to metabolic disease. The correlations were not enormous, and the cross-sectional design cannot establish direction of causation — it is possible that poor gut health disturbs sleep rather than the reverse, or that both are downstream of something else — but the pattern is consistent with the more mechanistic animal work, in which circadian-disrupted mice rapidly develop microbiome changes and metabolic impairment, and in which transferring the disrupted-mouse microbiome into a healthy mouse transfers some of the metabolic phenotype.

The link with the diurnal oscillation findings of chapter 20 is not coincidental. The gut microbiome itself oscillates over the day, with different taxa rising and falling in predictable patterns driven by the host's feeding schedule, host-derived circadian signals, and the local metabolic environment [REF:thaiss2014] [REF:thaiss2016]. Sleep disruption is fundamentally a disruption of the circadian clock, and it is the loss of *oscillation* — the loss of the orderly alternation between active and quiescent states — rather than any single static change that appears to be the proximate mechanism. The gut bacterium that was supposed to bloom at three in the morning and fade by six, in a well-entrained human, loses that rhythm when the host loses theirs, and the downstream metabolic outputs that depend on the rhythmic alternation are diminished. The analogy to the light section of

this chapter is deliberate: the gut responds not to averages but to patterns, and when the patterns are smeared, the ecosystem responds by losing structure.

Stress is, if anything, a harder subject to write about honestly, because the word “stress” covers so many physiologically different states — acute versus chronic, physical versus psychological, controllable versus uncontrollable, novel versus familiar — and because the human trials in this area are mostly small, short, and prone to confounding by the dietary and behavioural changes that stressed people make. The clearest animal evidence comes from the maternal-separation paradigm, in which rat or mouse pups are briefly separated from their mothers during a critical developmental window and develop lifelong alterations in gut microbiome composition and in stress-axis reactivity. The clearest human evidence is bidirectional: people with chronic anxiety and depression have, on average, gut microbiomes that differ from healthy controls, with reductions in several taxa that produce neuroactive metabolites; and, in the other direction, short-term psychological stressors (exam periods, combat deployment, spaceflight) produce measurable short-term shifts in gut community structure.

What is harder to say — and I want to be explicit about this — is what any given person should *do* about the stress-microbiome link, because the intervention literature is thin. There is some evidence that mindfulness-based interventions and cognitive behavioural therapy shift the gut microbiome in favourable directions in parallel with their mental health effects, but it is currently impossible to say whether the microbiome changes are *driving* the improvement, *responding* to it, or simply *accompanying* it because mood and eating habits shift together. For now, the honest summary is: there is a dialogue between the nervous system and the gut microbiome, the dialogue runs both ways, and interventions that reduce chronic stress also tend to move the microbiome in the direction we associate with better health. Whether any of this warrants calling meditation a “microbiome intervention” is a matter of taste; in strict evidential terms, the case is not yet made.

What I will say confidently is this: the ordinary human things — getting enough sleep, protecting the dark at night, allowing the body to rest and recover, having stable daily rhythms, maintaining social and family connection — are, on current evidence, good for the gut microbiome as well as for everything else we knew they were good for. The microbiome provides yet another reason to treat sleep and stress as first-class health inputs rather than as inconveniences to be worked around.

21.10 The Medicine Cabinet: Non-Antibiotic Drugs and the Microbiome

Antibiotics will get their own full chapter, because their effects on the gut microbiome are so profound and because their story is a story in itself. What I want to address in this section is something subtler and, for most readers, more relevant: the effects of the non-antibiotic medications that millions of people take every day, often for years at a time, without any idea that they are modifying their gut ecosystem in the process.

The paper that forced this question into the mainstream of microbiome research was a 2018 *Nature* study by Lisa Maier and colleagues at EMBL Heidelberg [REF:maier2018]. The researchers took the library of pharmacological compounds routinely used in human medicine — more than a thousand of them, spanning every major therapeutic category from blood pressure medications to antihistamines to antipsychotics — and screened each of them, one at a time, against a panel of forty representative human gut bacterial strains grown in the laboratory. They asked a simple question: at doses comparable to those achieved in the human gut during normal therapeutic use, how many

of these “non-antibiotic” drugs actually suppress the growth of human gut bacteria?

The answer, astonishingly, was nearly a quarter. Of the roughly 835 compounds that reached the gut at plausible concentrations, 24% inhibited the growth of at least one of the tested strains, and many inhibited several. The drugs that came out as strongly antimicrobial included proton-pump inhibitors (used for heartburn and acid reflux), certain antipsychotics (including chlorpromazine), some of the anti-diabetic medications, calcium channel blockers, and several classes of drug that had never been suspected of having any effect on bacteria at all. A proper antibiotic, the researchers noted, typically inhibits gut bacteria at far lower concentrations than these drugs do — but at the concentrations actually present in the gut of a person taking these medications long-term, the effects are real and measurable.

The implication is important. We have been dosing large fractions of the adult population for years or decades with drugs that quietly exert selection pressure on the gut microbiome, without anyone — not the prescribers, not the regulators, not the patients — having any framework for assessing what this cumulative pressure might be doing. The specific example that has become best characterised is proton-pump inhibitors (PPIs), a class of drugs that includes omeprazole, lansoprazole, esomeprazole, and pantoprazole and that is prescribed to hundreds of millions of people worldwide for heartburn, gastric ulcer prevention, and reflux disease.

PPIs suppress acid production in the stomach. One of the oldest functions of gastric acid is to sterilise incoming food and water — to kill off the microbes that enter the mouth and keep them from reaching the small intestine and colon alive. When you turn the acid off, oral and upper-respiratory bacteria that would normally have been killed in transit survive and establish themselves further down. A 2016 *Gut* paper by Floris Imhann and colleagues [REF:imhann2016] analysed stool microbiomes from over 1800 Dutch volunteers, some of whom were taking PPIs long-term, and found that PPI users had significantly altered gut microbial communities — specifically, increased abundance of oral bacteria (including *Streptococcus* and related genera) in the gut, increased abundance of *Enterococcus*, and shifts in several other markers consistent with upper-GI bacteria colonising the lower bowel. The effect size was substantial — comparable in magnitude to the effect of a course of antibiotics — and it was dose-dependent on the duration of PPI use.

Whether this shift causes clinical harm is a separate and still-evolving question. PPI users have been reported to have a modestly increased risk of *Clostridium difficile* infection, of small intestinal bacterial overgrowth, and possibly of community-acquired pneumonia, all of which are biologically consistent with loss of the acid barrier. The risks are not enormous for any individual, but they are real, and they add up when the drug is taken for years as many people do. I do not want this section to be read as telling every PPI user to stop their medication; PPIs are valuable drugs, and for patients with severe reflux disease or with high ulcer risk, the benefits clearly exceed the microbiome cost. What it should be read as telling is this: the assumption that a drug is “only” acting on its declared target, and that its effects outside that target can be ignored, is an assumption that microbiome research has comprehensively overturned. Whenever a non-antibiotic drug is taken for a prolonged period, it is worth asking — with a prescriber’s help — whether the duration is still necessary, whether the dose is still the minimum that achieves the goal, and whether alternative approaches might reduce the cumulative burden of chronic pharmacological pressure on the gut ecosystem.

Metformin, the first-line drug for type 2 diabetes, provides a particularly interesting counter-example. Metformin also alters the gut microbiome — a 2015 *Nature* paper by Karin Forslund and colleagues established this definitively, by showing that microbiome differences once attributed to

type 2 diabetes itself were in fact largely due to metformin treatment [REF:forslund2015]. What is interesting is that the microbiome changes metformin induces appear to be part of *how the drug works*: metformin increases the abundance of several short-chain-fatty-acid-producing bacteria and of *Akkermansia muciniphila*, and these shifts contribute to the drug's glucose-lowering effect. In germ-free mice, some of metformin's metabolic benefits are substantially attenuated, and when the "metformin-shifted" human microbiome is transferred to germ-free mice, the mice show improved glucose handling even without the drug. The microbiome-mediated component of metformin's effect is not a side effect; it is part of the mechanism. This is almost certainly not unique to metformin — other drugs are likely to work, at least partly, through their effects on the gut microbiome, and the field is only beginning to ask systematically which ones.

The general point I want to leave the reader with is that almost every oral medication meets the gut microbiome before it meets its target, and the encounter is rarely without consequence. If you are on a long-term medication and you have never asked how it interacts with your gut flora, the answer is probably "we do not fully know, but it almost certainly does something." For most drugs most of the time, the net balance is positive; the drug's benefits outweigh the microbiome costs. But the balance is not zero, and pretending that it is — as we did for most of the last century — is no longer intellectually honest.

21.11 The Water We Drink: Chlorine, Hardness, and the Microbes That Aren't There

I had not intended, when I started this chapter, to write a section on drinking water. It seemed too narrow a topic. But as I worked through the environmental inputs to the microbiome one after another, it became difficult to leave water out, because the argument about water parallels the argument about air so closely: we have taken a substrate that once carried a rich and seasonally varying microbial community and have processed it, for perfectly understandable historical reasons, into something nearly microbially inert. The question of whether that was a good trade-off on net is, once again, one that the microbiome literature is starting to ask.

Chlorination of public drinking water was one of the most successful public health interventions of the twentieth century. Before chlorination, waterborne cholera, typhoid, and dysentery killed millions every year in the industrialised world; after chlorination, these diseases essentially disappeared from places with treated water supplies. Nothing I write in this section should be read as questioning that history. Chlorine in the municipal water supply has saved more lives than almost any drug we have invented since, and it continues to save lives every day in countries where water treatment has recently been introduced. The microbiome question is not whether chlorination is good — it is, and it remains so — but whether the particular version of "clean water" we deliver to modern homes has microbial consequences beyond the absence of pathogens that deserve acknowledgement.

The honest answer is: probably, but we know less than we should. A handful of studies have examined whether long-term consumption of chlorinated tap water (versus bottled water, or filtered water, or well water) correlates with differences in the gut microbiome. The results are mixed and the studies are small; there is no convincing evidence that drinking chlorinated water at normal municipal concentrations measurably shifts the gut microbiome in a meaningful way. The chlorine is substantially consumed in transit through the digestive tract, is diluted by stomach contents, and is further neutralised by organic matter in food, such that the concentration that reaches the distal

small intestine and colon is low. If there is an effect, it is small and has not been convincingly isolated from the many other variables that distinguish tap-water drinkers from bottled-water drinkers.

More interesting, and better documented, is the effect of water *hardness* — the mineral content, particularly calcium and magnesium — on the gut microbiome. Several observational studies have linked residential water hardness with gut microbiome composition, with harder water (higher mineral content) correlating with modest differences in community structure and, in some studies, with differences in the abundance of specific short-chain fatty acid producers. The mechanistic picture is not yet clear; calcium and magnesium are known to interact with bile acid metabolism in ways that could plausibly feed back into microbial community structure. This is the kind of subtle environmental variable whose effects are probably real but whose magnitude is unlikely to matter much compared with diet or antibiotic exposure.

What does seem worth mentioning — and what actually prompted me to include this section — is the contrast between the water we drink and the water our ancestors drank. For most of human history, drinking water came from streams, rivers, springs, wells, and rain barrels, and it carried with it a living community of environmental microbes. Some of those microbes were dangerous, which is why chlorination was introduced; but many were harmless or tonic, part of the “Old Friends” stream of exposure that Graham Rook describes. When we replaced stream water with treated water, we did not just remove the pathogens; we removed the entire microbial content, pathogen and tonic friend alike. The calorie-free, pathogen-free, essentially sterile water that flows from a modern tap is an evolutionary novelty, and although no one is seriously proposing that we should return to drinking untreated river water, it is worth pausing to notice that the water input to the microbiome has been one more channel by which modern life has removed environmental microbial exposure from the daily experience of the human gut.

I would not make any drinking-water recommendations on the basis of current microbiome evidence. Municipal treated water is safe and should be drunk; bottled water has no advantage and a substantial environmental disadvantage; well water is fine if it has been properly tested. The point of this short section is only to note that water, too, belongs on the list of environmental inputs that modern sanitation has quietly shifted, and that the microbiome question will eventually need to be asked even for a substrate we think of as elementally simple.

21.12 Taking Stock: The Modern Environment as a Microbiome Experiment

If I have done my job in this chapter, the reader will by now be carrying a picture that is subtly but importantly different from the one most of us grew up with. The picture is this: a modern sanitised home, a climatized office, a filtered air supply, a paved neighbourhood, a pet-free apartment, a chlorinated tap, a bright screen at midnight, a mild chronic sleep debt, and a long-term maintenance medication — each of these, taken on its own, is a small environmental departure from the conditions under which the human microbiome evolved. Taken together, they are a substantial and cumulative one. They are not, individually, harmful in any dramatic way; nobody is about to die because their bedroom is too warm at night. But they add up, and the thing they are adding up to is a quiet, multi-decade, uncontrolled experiment on the human microbiome at population scale, an experiment for which none of us ever signed a consent form and whose long-term outcomes we are still learning to measure.

The honest scientific position is that we do not yet know how large the net effect of all these envir-

onmental departures is, or how much of the rise in allergic, autoimmune, metabolic, and possibly neuropsychiatric disease in industrialised populations it accounts for. The candidate mechanisms are plausible and increasingly well-documented; the direct causal evidence from controlled human trials is thin, and will remain thin for a long time, because you cannot randomise people to sleep in darkness for twenty years or to grow up with a dog versus without one. What we have instead is converging evidence from animal work, natural experiments (like the Amish-Hutterite comparison), observational cohorts, and short-term interventional trials, all of which point in the same direction without any one of them being definitive. A reasonable person can, and should, look at that evidence base and conclude that the cumulative environmental modifications of the last century are *probably* making some meaningful contribution to the modern disease pattern, while being careful not to overclaim about the magnitude.

What *should* a reader take away from this chapter in the form of practical changes they might actually make? I want to be specific, because the biggest risk of a chapter like this is that it leaves the reader feeling simultaneously alarmed and paralysed. In descending order of how well the evidence supports the recommendation, here is what I think the current literature warrants:

Sleep in the dark, and let the morning be bright. The evidence on circadian disruption and the microbiome is among the strongest in this chapter. Black-out curtains in the bedroom, no bright screens in the last hour before sleep, and deliberate morning light exposure together cost almost nothing, are supported by both mechanism and outcome data, and are beneficial for reasons that go well beyond the microbiome.

Open the window every day, in almost any weather. The HEPA and household air evidence suggests that microbial variety, not microbial purity, is what matters for a healthy home, and the simplest and cheapest way to deliver variety is to let the outside in. Fifteen minutes a day is enough.

If you are considering a dog, and especially if you have young children, bring the dog in early. The pet-microbiome evidence is unusually clean for this field, and the effect is substantial enough to partially compensate for the microbiome disruptions caused by caesarean delivery and early antibiotic exposure.

Let young children get dirty. Soil, grass, gardens, small farms, puddles, mud — these are not hazards to be managed but inputs to be welcomed. The Old Friends evidence, while theoretical in places, is the best explanation we currently have for the epidemic of allergy and autoimmune disease in industrialised populations, and the practical response to it is simply to stop preventing the ordinary childhood contact with environmental microbial life that used to be taken for granted.

Move the body, preferably outside. The exercise-microbiome link is real, the outdoor-exercise link is probably real for a second and additional reason (increased environmental exposure), and the dose need not be extreme. Thirty minutes of walking in a park most days delivers much of what the literature currently supports.

Let the house experience some seasons. The thermostat section is the most speculative part of this chapter, and I acknowledged that it was. But the general idea that a body and its microbiomes *expect* some variation in temperature and humidity over the year, and that holding everything constant is a novel experiment whose consequences are incompletely understood, seems worth taking seriously enough to let the house be a degree or two warmer in summer and cooler in winter than you might otherwise default to.

Be mindful of long-term medication. Not fearful; mindful. If you have been on a proton-pump inhibitor for five years without ever revisiting whether the indication still holds, that is worth a

conversation with your doctor — not because the drug is dangerous, but because the principle of minimum necessary duration is a good one when a medication is interacting with the gut microbiome every day.

What this chapter has *not* covered is the intervention most likely to be the single largest environmental disruptor of the human microbiome in the modern era: antibiotics. I have deferred antibiotics deliberately, because their effects are large enough and their story is complicated enough that they deserve their own chapter, not a subsection in this one. Chapter 22 takes up that story, and it is where much of what we have discussed in this chapter and the previous one will be brought into contact with the most powerful tool modern medicine has for reshaping microbial communities — for good and for ill.

Key References

- [REF:wu2018] Wu G, Tang W, He Y, et al. Light exposure influences the diurnal oscillation of gut microbiota in mice. *Biochemical and Biophysical Research Communications*. 2018;501(1):16-23. doi:10.1016/j.bbrc.2018.04.095
- [REF:lee2022] Lee CC, Liang F, Lee IC, Lu TH, Shan YY, Jeng CF, Zou YF, Yu HT, Chen Alen SK. External light-dark cycle shapes gut microbiota through intrinsically photosensitive retinal ganglion cells. *EMBO Reports*. 2022;23(6):e52316. doi:10.15252/embr.202052316
- [REF:brooks2021] Brooks JF, Behrendt CL, Ruhn KA, et al. The microbiota coordinates diurnal rhythms in innate immunity with the circadian clock. *Cell*. 2021;184(16):4154-4167.e12. doi:10.1016/j.cell.2021.07.001
- [REF:rogers2021] Rogers AE, Hu YJ, Yue Y, Wissel EF, Petit RA 3rd, Jarrett S, Christie J, Read TD. Shiftwork, functional bowel symptoms, and the microbiome. *PeerJ*. 2021;9:e11406. doi:10.7717/peerj.11406
- [REF:lopezsantamarina2023] Lopez-Santamarina A, Mondragon ADC, Cardelle-Cobas A, Santos EM, Porto-Arias JJ, Cepeda A, Miranda JM. Effects of unconventional work and shift work on the human gut microbiota and the potential of probiotics to restore dysbiosis. *Nutrients*. 2023;15(13):3070. doi:10.3390/nu15133070
- [REF:hong2020] Hong F, Pan S, Xu P, Xue T, Wang J, Guo Y, Jia L, Qiao X, Li L, Zhai Y. Melatonin orchestrates lipid homeostasis through the hepatointestinal circadian clock and microbiota during constant light exposure. *Cells*. 2020;9(2):489. doi:10.3390/cells9020489
- [REF:thaiss2014] Thaiss CA, Zeevi D, Levy M, et al. Transkingdom control of microbiota diurnal oscillations promotes metabolic homeostasis. *Cell*. 2014;159(3):514-529. doi:10.1016/j.cell.2014.09.048
- [REF:thaiss2016] Thaiss CA, Levy M, Korem T, et al. Microbiota diurnal rhythmicity programs host transcriptome oscillations. *Cell*. 2016;167(6):1495-1510.e12. doi:10.1016/j.cell.2016.11.003
- [REF:krisko2020] Krisko TI, Nicholls HT, Bare CJ, et al. Dissociation of adaptive thermogenesis from glucose homeostasis in microbiome-deficient mice. *Cell Metabolism*. 2020;31(3):592-604.e9. doi:10.1016/j.cmet.2020.01.012
- [REF:richardson2019] Richardson M, Gottel N, Gilbert JA, Lax S. Microbial similarity between students in a common dormitory environment reveals the forensic potential of individual microbial

signatures. *mBio*. 2019;10(4):e01054-19. doi:10.1128/mBio.01054-19

[REF:maestre2018] Maestre JP, Jennings W, Wylie D, Horner SD, Siegel J, Kinney KA. Filter forensics: microbiota recovery from residential HVAC filters. *Microbiome*. 2018;6(1):22. doi:10.1186/s40168-018-0407-6

[REF:stein2016] Stein MM, Hrusch CL, Gozdz J, et al. Innate immunity and asthma risk in Amish and Hutterite farm children. *New England Journal of Medicine*. 2016;375(5):411-421. doi:10.1056/NEJMoa1508749

[REF:rook2013] Rook GAW. Regulation of the immune system by biodiversity from the natural environment: an ecosystem service essential to health. *Proceedings of the National Academy of Sciences*. 2013;110(46):18360-18367. doi:10.1073/pnas.1313731110

[REF:rook2014] Rook GAW, Raison CL, Lowry CA. Microbial “Old Friends”, immunoregulation and socioeconomic status. *Clinical and Experimental Immunology*. 2014;177(1):1-12. doi:10.1111/cei.12269

[REF:tun2017] Tun HM, Konya T, Takaro TK, et al. Exposure to household furry pets influences the gut microbiota of infants at 3-4 months following various birth scenarios. *Microbiome*. 2017;5(1):40. doi:10.1186/s40168-017-0254-x

[REF:nielsen2020] Nielsen CC, Gascon M, Osornio-Vargas AR, et al. Natural environments in the urban context and gut microbiota in infants. *Environment International*. 2020;142:105881. doi:10.1016/j.envint.2020.105881

[REF:lax2014] Lax S, Smith DP, Hampton-Marcell J, et al. Longitudinal analysis of microbial interaction between humans and the indoor environment. *Science*. 2014;345(6200):1048-1052. doi:10.1126/science.1254529

[REF:vallescolomer2023] Valles-Colomer M, Blanco-Míguez A, Manghi P, et al. The person-to-person transmission landscape of the gut and oral microbiomes. *Nature*. 2023;614(7946):125-135. doi:10.1038/s41586-022-05620-1

[REF:clarke2014] Clarke SF, Murphy EF, O’Sullivan O, et al. Exercise and associated dietary extremes impact on gut microbial diversity. *Gut*. 2014;63(12):1913-1920. doi:10.1136/gutjnl-2013-306541

[REF:motiani2020] Motiani KK, Collado MC, Eskelinen JJ, et al. Exercise training modulates gut microbiota profile and improves endotoxemia. *Medicine and Science in Sports and Exercise*. 2020;52(1):94-104. doi:10.1249/MSS.0000000000002112

[REF:benedict2016] Benedict C, Vogel H, Jonas W, et al. Gut microbiota and glucometabolic alterations in response to recurrent partial sleep deprivation in normal-weight young individuals. *Molecular Metabolism*. 2016;5(12):1175-1186. doi:10.1016/j.molmet.2016.10.003

[REF:sun2023sleep] Sun J, Fang D, Wang Z, Liu Y. Sleep deprivation and gut microbiota dysbiosis: current understandings and implications. *International Journal of Molecular Sciences*. 2023;24(11):9603. doi:10.3390/ijms24119603

[REF:maier2018] Maier L, Pruteanu M, Kuhn M, et al. Extensive impact of non-antibiotic drugs on human gut bacteria. *Nature*. 2018;555(7698):623-628. doi:10.1038/nature25979

[REF:imhann2016] Imhann F, Bonder MJ, Vich Vila A, et al. Proton pump inhibitors affect the gut microbiome. *Gut*. 2016;65(5):740-748. doi:10.1136/gutjnl-2015-310376

[REF:forslund2015] Forslund K, Hildebrand F, Nielsen T, et al. Disentangling type 2 diabetes and metformin treatment signatures in the human gut microbiota. *Nature*. 2015;528(7581):262-266. doi:10.1038/nature15766

End of Chapter 21 draft.

Chapter 22: Drugs Beyond Antibiotics — The Quiet Pharmacology of the Gut

22.1 The Quietest Side Effect

In 2018, a small group of researchers at the European Molecular Biology Laboratory in Heidelberg did something that, in retrospect, ought to have been done thirty years earlier and was not. Lisa Maier, Mihaela Pruteanu, Michael Kuhn, and their colleagues took the routine pharmacological library that any university medicinal-chemistry department keeps in its freezer — more than a thousand drugs in regular human medical use, spanning every major therapeutic category from blood-pressure pills to antihistamines to antipsychotics to anti-cancer agents — and exposed each one, individually, to a panel of forty representative human gut bacteria grown in the laboratory. They asked an embarrassingly simple question: at concentrations that the gut of an actual patient would experience during normal therapeutic use, how many of these supposedly “non-antibiotic” drugs actually inhibit the growth of the human gut microbiota? [REF:maier2018]

The answer, when their paper appeared in *Nature*, was nearly a quarter. Of the roughly 835 compounds that reached the gut at plausible concentrations, 24 percent inhibited the growth of at least one of the tested strains, and many inhibited several. The list of culprits made for uneasy reading. Proton pump inhibitors — the drugs hundreds of millions of people swallow every day for reflux. Several major antipsychotics, including chlorpromazine. Calcium-channel blockers used for blood pressure. Anti-diabetic agents. A handful of common painkillers. None of these had ever been designed to interfere with bacteria. None had ever been licensed with that intention. None were called “antibiotics” by any regulatory authority on Earth. And yet in the dish, at human-relevant concentrations, they behaved as though they were.

I want to start the chapter with this paper because it overturned a comfortable assumption that had quietly governed medicine for most of the twentieth century. The assumption was this: a drug acts on its declared target. A blood-pressure pill acts on blood pressure. An antihistamine acts on the H1 receptor in the lining of your nose. An antipsychotic acts on dopamine receptors in the brain. Whatever else the drug is doing — whatever it incidentally bumps into on its way through the gut — is either negligible, or shows up as a side effect noticed by patients and tabulated in the package insert. The Heidelberg paper made this assumption untenable. Most oral medications meet the gut microbiome before they meet anything else, and the encounter, the data showed, is rarely without consequence for the bacteria. The bacteria simply have not been complaining loudly enough for us to notice.

In Chapter 21, I introduced this idea briefly, in section 21.10, as one of several “other levers” that shape the modern microbiome. I told you about Maier’s screen, said a few words about proton pump inhibitors and metformin, and promised that the topic deserved its own chapter. This is that chapter. What I want to do here is take the question seriously: across the full range of drugs that ordinary people take across an ordinary lifetime, which ones substantially shift the gut ecosystem, by how much, and for how long, and — crucially — does it matter?

I will tell you in advance what kind of chapter this is and what kind of chapter it is not. It is not

a list of drugs to avoid. I am an emergency physician; I have spent most of my professional life seeing patients whose lives have been saved or made bearable by exactly the medications discussed in the pages that follow. Telling people to throw away their proton pump inhibitor because it shifts their *Lactobacillus* count would be irresponsible to the point of dangerous, and I will not be doing it. The chapter is also not a guide to “alternatives”. The question of whether an existing medication should be changed, paused, deprescribed, or replaced is a question for the patient and the doctor who knows the patient — not for a popular-science book that knows neither. What this chapter is, and what I hope it manages to be, is an honest map of where the evidence currently stands, so that readers who are taking long-term medication know what their drug is doing inside them, can have an informed conversation with their treating doctor about it, and can make sensible adjustments — to diet, to fibre intake, to fermented foods — that are within their own power and that the evidence reasonably supports.

There is one more thing I want to set out at the beginning. Most of the drugs in this chapter have positive reasons to exist. They prevent strokes, they prevent heart attacks, they keep people alive who would otherwise be dead. Their effect on the microbiome is, in almost every case, a *cost* set against a much larger *benefit*. The arithmetic, on the available evidence, comes out in favour of the benefit. What the microbiome story changes is not whether to take the drug; it is the hidden ledger on which the taking is recorded. Knowing the cost is on the ledger allows the cost to be discussed, monitored, and — where necessary, and only in conversation with a clinician who knows the case — minimised. That is the only practical use to which this chapter should be put.

22.2 Reading the Evidence: A Note Before We Start

Before I start describing individual drug classes, I want to spend a paragraph or two on the evidence itself, because the evidence in this field is uneven in a way the reader should know about.

For some drugs the evidence is, by the modest standards of microbiome research, very strong. For metformin we have multiple large human cohorts, a meta-analysis pooling more than a thousand individuals, mechanistic experiments in germ-free mice, and a clear molecular story about how the drug talks to specific bacterial taxa. For proton pump inhibitors we have several large observational cohorts and a coherent biological mechanism — remove the gastric acid barrier, and microbes that would normally have been killed in the stomach reach the small bowel in living form. For these drug classes, the science has reached the point where I can write declarative sentences without needing to hedge every other word.

For other drugs the evidence is much thinner. For ordinary H1 antihistamines — the cetirizines and loratadines that millions of hay-fever sufferers swallow every spring — the published microbiome literature is, as far as I can find, almost nothing. For chronic laxative use we have one really good landmark paper from 2018 [REF:tropini2018] and very little since. For opioids we have rich and convincing evidence from mouse experiments and almost no proper human cohort data. In each of these cases I am going to tell you what the mechanistic story looks like, but I am also going to tell you that the human evidence has not been done, or has been done badly, or has been done in groups too small to draw conclusions from. Where I am extrapolating from animals or from theory, I will say so. Where I am guessing, I will say that too.

Finally, the field has a recurring problem with what statisticians call confounding by indication. People who take a drug for years are different from people who do not, in many ways other than the

drug. Someone on long-term proton pump inhibitors usually has reflux or has been on aspirin or both, and reflux and aspirin are themselves things that can shift the gut community. Someone on long-term antipsychotics usually has schizophrenia or bipolar disorder, and severe mental illness is itself associated with a different gut microbiome — it is not clear which way the arrow points. The cleanest evidence comes from studies that follow the same person before and after starting a drug, or from randomised trials, or from animal experiments in which the only variable is the drug. Where I am drawing on those, the evidence is more trustworthy than where I am drawing on a one-time snapshot of a clinic population. I will not stop the narrative every paragraph to drill this in, but the reader should keep it in the back of their mind throughout.

With those caveats acknowledged, let me begin where the gut begins — with the acid in the stomach.

22.3 Closing the Acid Gate: PPIs and Their Cousins

A wall that is almost always there

The human stomach is, for most of every day, an acid bath. Its pH sits between 1.5 and 3 — about the acidity of lemon juice, sometimes more — and it is held there by the parietal cells of the stomach lining, which actively pump hydrogen ions into the lumen via an enzyme called the gastric proton pump. The acid is there for three reasons. It denatures dietary proteins so digestive enzymes can attack them. It activates pepsin, the stomach's main protein-cleaving enzyme. And — most relevant for this chapter — it kills almost everything that comes in through the mouth.

Drink a glass of water, swallow a mouthful of food, take a deep breath of air with a few aerosolised bacteria in it, and the great majority of the live microbial cells you ingest will be dead within minutes. The acid wall is the front line of what immunologists once called *colonisation resistance* and what an ecologist would simply call a hostile climate at the door of the ecosystem. It is the reason you can eat a slightly off sandwich and not get violently ill. It is the reason cholera has to evolve specific acid-resistance machinery to survive long enough to reach the gut and infect you. It is the reason the small bowel of a healthy adult contains a few hundred bacteria per millilitre of its contents, while the colon a metre further down contains a hundred billion. The acid does most of the sorting.

Now imagine what happens if you switch the wall off.

A *proton pump inhibitor*, or PPI, is a drug that turns the wall down — though, importantly, not permanently. The molecule binds tightly and covalently to an individual proton pump and disables that particular pump molecule for good. But the pump is not a permanent fixture of the parietal cell. It is a protein, and like almost all proteins it is constantly being replaced. The stomach is a harsh environment even for the cells that make the acid, and the pump itself wears out: its own half-life in the membrane is on the order of two days, and the parietal cell replaces it on a rolling basis regardless of whether the drug is present or not. When the next dose of the PPI does not arrive, new pumps keep being made, old (drug-bound) pumps keep being degraded, and within a day or two the stomach is back to its normal acid-secreting self. Recovery of acid secretion after a dose of omeprazole happens with a half-time of roughly twenty-eight hours in humans, and a little quicker for lansoprazole [REF:shin2002ppi]. For most PPIs, an additional mechanism helps along the return to baseline: the cell's own glutathione — the same antioxidant that handles day-to-day molecular housekeeping — can partly reverse the drug's binding to the pump, releasing the inhibitor again. So although each individual pump–drug interaction is covalent and irreversible at the molecular level, the effect on the patient's stomach is very much reversible, and the acid wall

is rebuilt from scratch every couple of days. This matters because it tells us something reassuring: stopping a PPI does not leave a permanently disabled stomach behind. The machinery comes back. The first PPI, omeprazole, was approved in the late 1980s; the class now includes lansoprazole, esomeprazole, pantoprazole, and rabeprazole, and is one of the most-prescribed drug classes in the world. Estimates vary by country, but in most of the developed world somewhere between five and ten percent of the adult population is on a PPI on any given day, often for years at a stretch, sometimes for decades. They are excellent drugs for what they were designed to do: they heal peptic ulcers, they treat severe reflux disease, and they prevent stomach bleeding in people taking long-term aspirin or anti-inflammatories. For a patient with bad oesophagitis or a recent gastric bleed, a PPI is, without exaggeration, the difference between misery and a normal life. Nothing in this section should be read as suggesting otherwise.

What the PPI does that no one really thought about for the first twenty years of the drug's existence is open the front door.

What gets through

Frank Imhann and colleagues, working in the Netherlands, published the first really convincing human-cohort paper on this question in *Gut* in 2016 [REF:imhann2016]. They sequenced the gut bacteria of nearly two thousand people in the LifeLines DEEP cohort, a prospective study of the population of the northern Netherlands, and asked whether PPI users had different microbiomes from non-users. They did. The differences were not subtle. Bacteria that normally live in the mouth and upper throat — *Streptococcus*, *Veillonella*, *Rothia*, *Actinomyces*, and various other oral commensals — were substantially more abundant in the lower gut of PPI users. The reason was straightforward. With the acid wall down, mouth bacteria that are continuously swallowed in saliva (we swallow about one to two litres a day) were no longer being killed in the stomach. They survived the trip. They reached the small bowel and the colon, where, in many cases, they began to colonise tissue they had never reached before.

The Imhann team also noted that the magnitude of these shifts, when adjusted for confounders, was *larger* than the shifts observed in the same dataset between people with and without irritable bowel syndrome, or with and without inflammatory bowel disease. That is a striking comparison. Taking a daily pill of omeprazole was, by their analysis, a bigger statistical perturbation of the gut community than having an inflammatory bowel disease.

Independent work has consistently found the same thing. Subsequent studies in larger cohorts and in different populations have reproduced the core finding — PPI users have an oralised gut microbiome — and have begun to add detail to the mechanism. One recent line of work has shown that PPIs also increase the permeability of the intestinal lining itself, by reducing the expression of tight-junction proteins (occludin in particular) that hold gut epithelial cells together [REF:zhang2025ppi]. A leakier gut is a more inflamed gut, and a more inflamed gut is one in which previously commensal microbes can begin to cause trouble.

The clinical consequences

What does any of this mean for an actual patient on a PPI?

The honest answer is: in most cases, very little that they would notice in any given week. The microbiome shifts are real and measurable, but the corresponding clinical risks — although they exist — are individually small in absolute terms, even if they are statistically significant in large

populations. Long-term PPI users have a modestly increased risk of *Clostridioides difficile* infection (the bug that causes the most feared form of post-antibiotic diarrhoea), of small intestinal bacterial overgrowth (SIBO, in which the small bowel ends up containing a colon-like density of bacteria it is not equipped to handle), and probably of community-acquired pneumonia from inhaled oral microbes that are no longer being cleared from the throat by adequate gastric backstop. For most healthy adults who happen to be taking a PPI for two months for a peptic ulcer, none of these risks ever materialises. For an eighty-year-old in a nursing home who has been on a PPI for fifteen years and who takes it because somebody put her on it for vague indigestion in 2010 and never reviewed the prescription, the picture is different. The cumulative risk has had time to accumulate. The benefit, in many such cases, was never very large to begin with.

A 2025 meta-analysis of cancer immunotherapy outcomes added another wrinkle [REF:xu2025ppi]. Pooling sixty-nine studies and more than twenty-two thousand patients, the authors showed that taking a PPI within three months of starting an immune checkpoint inhibitor — the class of cancer drugs that includes pembrolizumab and nivolumab and that has revolutionised the treatment of melanoma, lung cancer, and several other malignancies — was associated with significantly worse overall survival, progression-free survival, and response rate. The mechanism is not certain, but the leading hypothesis involves the gut microbiome: the PPI shifts the community in a way that interferes with the systemic immune responses on which checkpoint inhibitor therapy depends. This is a striking result and one that, if it holds up, has direct implications for prescribing in oncology.

I want to be careful here. The Xu meta-analysis is an observational pooling, not a randomised trial; people who happen to be on PPIs while starting cancer immunotherapy are not a random subset of cancer patients, and the confounding-by-indication problem is alive and well. But the signal has appeared in enough independent cohorts that the major cancer centres have begun to discuss minimising PPI use around the start of checkpoint inhibitor therapy when this can be done safely. It is a good example of how the microbiome conversation, once it gets specific enough, starts to change real prescribing.

What “rational PPI use” looks like

The professional literature on PPI prescribing has slowly converged on a position that is worth sharing in plain language. PPIs are excellent drugs for the conditions they were designed to treat: severe oesophagitis, peptic ulcers, prevention of bleeding in patients on chronic NSAIDs or anti-platelet therapy, and a few other specific indications. For these conditions, the benefits clearly outweigh the microbiome cost. PPIs are *not* excellent drugs for vague indigestion, occasional heartburn, or “just in case” prescription continued long after the original reason for starting it has gone away. For these uses, the benefits are smaller and the cost — to the gut ecosystem, and very probably to bone density and to magnesium status as well, although those mechanisms are different — adds up over years.

If you have been on a PPI for more than a year and you have not had a conversation with your doctor about whether the indication still holds, that conversation is worth having [REF:sharma2025ppi]. It is not a conversation about stopping the drug. It is a conversation about whether the dose can be the lowest effective one, whether step-down to an as-needed pattern is reasonable, and whether the original reason still applies. For some patients the answer will be “yes, keep going at full dose” — that is fine, and the microbiome cost is an accepted part of the package. For others the answer will be “we can taper”, and the gut ecosystem will quietly thank them for it.

I should mention briefly the cousin class of acid-reducing drugs, the *H2 receptor antagonists* — ranitidine, famotidine, cimetidine — that do the same job as PPIs by a different molecular route (blocking the histamine receptor on the parietal cell rather than the proton pump itself). These drugs raise gastric pH less profoundly than PPIs do, and on theoretical grounds should have a smaller effect on the gut microbiome through the same acid-gate mechanism. The published evidence on this is, frankly, sparse — I could find very little serious recent work on H2 blockers and the gut microbiome at the time of writing — but the mechanistic prediction is that any microbiome effects should be in the same direction as PPIs but smaller in magnitude. For a patient who needs acid suppression and for whom a PPI feels excessive, an H2 blocker is a reasonable option, and the absence of dedicated microbiome data should not be read as evidence that the class is harmful in a way the class above is not.

22.4 The Drug That Works Through Microbes: Metformin

The puzzle of the missing mechanism

Metformin is the most-prescribed drug in the world for type 2 diabetes. It has been on the market since 1957 in France, since 1995 in the United States, and is taken by tens of millions of people on every continent. It is cheap, it is effective at lowering blood sugar, it modestly reduces cardiovascular events and possibly all-cause mortality, and — this is the embarrassing part — for most of its history nobody could explain exactly how it worked. Textbooks dutifully cited “activation of AMP-activated protein kinase” and “inhibition of hepatic gluconeogenesis”, but the molecular details kept refusing to add up. The drug clearly worked. The question of *why* it worked was strangely unsettled for a medicine taken by hundreds of millions of people.

In 2015, Karin Forslund and colleagues at the European Molecular Biology Laboratory in Heidelberg published a paper in *Nature* that began to provide an answer in a place no one had been seriously looking [REF:forslund2015]. The Forslund team did something subtle and important. They took a large dataset of 784 gut metagenomes from people with and without type 2 diabetes across three countries, and instead of asking the usual question — “how does the diabetic gut microbiome differ from the non-diabetic one?” — they asked a sharper one: “what happens when we separate diabetic patients who are on metformin from those who are not?”

The answer rearranged the picture. Untreated diabetics genuinely did have a distinctive and unhealthy-looking gut microbiome: they showed a real depletion of butyrate-producing genera such as *Roseburia* and *Subdoligranulum*, and a cluster of butyrate-producing *Clostridiales* — the very bacteria that nourish the gut lining and tamp down inflammation. That was the true disease signature, and it was consistent across the Danish, Swedish, and Chinese cohorts in the study. But previous papers had not separated treated from untreated patients, so the published “diabetic microbiome” was actually a muddle of two different signals pulling in different directions: the disease signature (depleted SCFA producers) and the drug signature (metformin partly reversing some of those depletions while adding effects of its own, such as an increase in *Escherichia* species and a decrease in *Intestinibacter*). When the Forslund team disentangled the two, the metformin-treated diabetics looked, in several respects, *more* like healthy controls than like their untreated counterparts — at least for the butyrate-producer taxa that the disease had stripped away. The drug was not causing the dysbiosis. It was partly repairing it, even as it introduced its own side-signatures along the way. This was a quietly devastating paper for several earlier microbiome–disease association studies that had not controlled for medication, and it was also

the moment when the question “does metformin work *through* the microbiome?” stopped being a fringe speculation and became a serious mechanistic hypothesis.

What metformin does to the gut community

The shifts that metformin produces in the gut bacterial community are now well characterised. A 2025 individual-patient meta-analysis pooling nine datasets and just over fourteen hundred type 2 diabetes patients [REF:su2025metformin] confirmed and refined the picture. Metformin reliably increases the abundance of two groups of bacteria. The first is short-chain-fatty-acid producers — bacteria that ferment dietary fibre into the small molecules (acetate, butyrate, propionate) that nourish the gut lining and modulate host metabolism. The second, more dramatically, is a single species: *Akkermansia muciniphila*, a comma-shaped anaerobe that lives in the mucus layer of the colon and has, over the last decade, become one of the most discussed organisms in metabolic microbiome research.

Akkermansia muciniphila is worth meeting properly, because it is going to come back several times in this chapter. Discovered by Muriel Derrien and Willem de Vos in 2004, it makes its living by eating the mucus your own intestinal cells secrete, breaking it down into nutrients it shares with neighbouring bacteria. In low-grade and well-regulated abundance, it appears to be protective — it is associated with leaner body weight, better glucose tolerance, lower systemic inflammation, and a less leaky gut barrier. In a healthy adult Western gut, *Akkermansia* might make up between one and four percent of the total bacterial cells; in obese and diabetic individuals, it is often substantially depleted; in people who have done well on metformin, it has typically rebounded toward and beyond healthy levels. The same species, somewhat surprisingly, has shown up as a positive predictor in studies of cancer immunotherapy response, of dietary intervention success, and of recovery from antibiotic damage. Whatever it is doing, the body seems to like having it.

The metformin–*Akkermansia* link is now firm enough that several research groups have begun the obvious experiment: give people *Akkermansia* directly, as an oral supplement, and see what happens. A 2024 randomised pilot trial gave a strain called *Akkermansia muciniphila* L2-7 to twenty-five men with type 2 diabetes who were already on stable metformin, and reported improved glycaemic variability and a small reduction in mean arterial blood pressure over fourteen days [REF:attaye2024]. The trial was small and the endpoints modest, but the principle held: a probiotic version of the species the drug is recruiting can produce some of the benefits the drug produces.

Is the microbiome the mechanism?

This is the question the literature is still arguing about, and the honest answer is: partly, yes, but not entirely.

The “yes” comes from several lines of evidence. When metformin is given to germ-free mice, some of its glucose-lowering effect is substantially attenuated. When the gut microbiota of metformin-treated humans is transferred into germ-free mice, the recipient mice show improved glucose tolerance even without the drug. When metformin is given to mice in which a key bile acid receptor (FXR) has been knocked out, much of the metabolic benefit disappears, suggesting that the drug works partly by changing the bile acids that microbes release — which then signal back to the host. A 2024 paper from the Xie group went further, showing that in a stress-induced mouse depression model, metformin reversed depressive-like behaviour through a chain that ran metformin → *Akkermansia muciniphila* enrichment → altered bile acids → reprogrammed tryptophan metabolism →

restored serotonin signalling [REF:xie2024metformin]. Faecal transfer from metformin-treated mice rescued the depressive behaviour in untreated mice. The whole chain is remarkable: a diabetes drug is producing an antidepressant effect through the gut microbiota.

The “not entirely” comes from the fact that metformin also has direct molecular effects on host cells. It does activate AMP-activated protein kinase. It does inhibit a particular complex of mitochondrial enzymes in liver cells. It does, by several routes, reduce the amount of glucose the liver dumps into the blood. A clever 2024 cohort analysis of metformin-treated diabetic men correlated the *Akkermansia/Romboutsia* ratio in the gut with cognitive performance scores, and with several plasma amino acid metabolites — suggesting that the metformin–microbiome–metabolome axis extends well beyond glucose control [REF:rosell2024]. But the same paper reminds the reader that not every metformin user gets every benefit, and the variability looks at least partly to be a function of who they had in their gut to begin with.

The current best summary, I think, is this. Metformin is a drug whose effects on the body are mediated by *both* a direct biochemical action on host cells and an indirect microbiome-mediated cascade, and the relative contribution of the two routes depends on the patient, on the dose, and on what was in their gut at the start of treatment. This is not the way the textbook chapters on metformin were written in the 1990s, when nobody knew there were bacteria involved. It is, I suspect, going to turn out to be the way many other drugs work as well — a host effect plus a bacterial effect, intertwined, the two often in the same direction but not always — and the microbiome story of metformin is the canonical example of what the rest of the field is starting to look like.

The unsettling implication

Let me close this section with a thought that I find more interesting than the science alone. If metformin works partly through *Akkermansia*, and *Akkermansia* abundance varies enormously between individuals (and across populations, and across diets), then the question of whether metformin will work in *you* is not solely a function of your liver or your insulin receptors. It is also a function of which microbes you happened to inherit from your mother, which ones you maintained or lost across decades of dietary choices, which antibiotics you took as a child, and what your gut community looks like on the day your doctor writes the prescription. Two patients with apparently identical type 2 diabetes can respond to metformin very differently, and a substantial fraction of that difference is now believed to be microbiome-driven. The clinical implication — still emerging, still being worked out — is that the same drug, given for the same indication, may benefit two patients to very different degrees because their gut ecosystems are not the same. Personalised medicine is often discussed as a matter of human genetics. The metformin story is a reminder that for at least some drugs, the microbial genome of the gut may matter as much as the human genome of the patient.

22.5 The Quiet Erosion: NSAIDs and the Small Bowel You Cannot Feel

A pain reliever that hurts where you cannot feel it

The non-steroidal anti-inflammatory drugs — NSAIDs for short — are among the most widely consumed pharmaceutical compounds on Earth. They include ibuprofen, naproxen, diclofenac, indomethacin, ketoprofen, mefenamic acid, the older aspirin, and the newer COX-2 selective inhibitors such as celecoxib. They are sold over the counter in most countries. They are taken for headaches,

period pain, sports injuries, low-grade fever, sprains, post-operative recovery, and chronic conditions such as osteoarthritis and rheumatoid disease. In the United States alone, an estimated seventy million prescriptions are written every year, and over-the-counter use is several times higher. Most users do not think of them as serious drugs at all.

Most people who think of NSAID side effects think of the stomach. NSAIDs reduce the protective mucus layer of the stomach lining by inhibiting the prostaglandin pathway that maintains it, and the resulting ulcers and bleeds — gastritis, gastric ulcer, duodenal ulcer, occasionally fatal upper gastrointestinal haemorrhage — are a well-known and well-publicised harm of long-term use. This is why doctors so often co-prescribe a PPI alongside long-term NSAIDs, and why aspirin is dispensed with that warning written on the bottle.

What is much less appreciated is that the lower part of the small bowel — the jejunum and the ileum, the section the doctor cannot easily reach with an endoscope and that the patient cannot easily feel — is also damaged by NSAIDs, and the damage there is, to a substantial extent, mediated by the gut microbiome.

What the small bowel goes through

The mechanism, worked out gradually across the last fifteen years, looks roughly like this. NSAIDs are absorbed in the small intestine, distributed via the bloodstream, and then secreted back into the bile. The bile carries the drug back into the duodenum, from where it travels down the small bowel a second time. Some NSAIDs (indomethacin, in particular, but also diclofenac and others) undergo this *enterohepatic recirculation* multiple times, exposing the small bowel mucosa to repeated waves of high local drug concentration that the systemic blood level does not reflect. At the same time, the drug damages the mitochondria of intestinal epithelial cells, weakens tight junctions between them, and reduces local prostaglandin-mediated repair.

So far this is purely a host-toxicity story. The microbiome enters the picture in the next step. As the epithelial barrier loosens, gut bacteria — or, more precisely, fragments of bacteria such as lipopolysaccharide from the outer membrane of Gram-negative organisms — translocate into the lining and trigger an inflammatory response through the TLR4/NF- κ B signalling pathway, the same innate-immunity pathway we met in the chapters on the immune system. The inflammation in turn damages more epithelium, the barrier loosens further, and a self-amplifying cycle of NSAID enteropathy is set in motion. Animal experiments have shown convincingly that the same dose of indomethacin produces dramatically more small-bowel damage in conventional mice than in germ-free ones. The drug is doing some of the damage. The bacteria are doing the rest.

A 2025 mouse study put numbers on this. Indomethacin-induced enteropathy was associated with characteristic dysbiosis: a loss of protective tight-junction proteins (ZO-1, claudin, occludin) and a shift in bacterial composition with an increase in *Firmicutes* and a loss of certain *Bacteroidetes* taxa. Restoring the barrier with a fermented food intervention reduced both the dysbiosis and the diarrhoea rate by about fifty percent [REF:varela2025]. A separate experimental line has shown similar findings with high-altitude hypoxia plus aspirin — a model relevant to climbers and to soldiers — in which the dysbiotic and inflammatory damage was substantially worse than with either insult alone, and was attenuated by a plant-derived compound that protected the intestinal barrier and rebalanced the microbial community [REF:tuerxuntayi2025].

In humans the effect is harder to study, because most of the small bowel is essentially inaccessible to ordinary clinical investigation. But studies that have looked — using video capsule endoscopy to

image the otherwise unreachable middle of the small bowel — have repeatedly shown small ulcers and erosions in a substantial fraction of long-term NSAID users who have no symptoms whatever and no abnormalities in routine blood tests. The injury is real, it is widespread, and most of the people carrying it have no idea.

Aspirin is a special case

I want to single out aspirin for a separate mention, because aspirin has a strange and double-edged relationship with the gut microbiome.

Low-dose aspirin (typically 75 to 100 mg daily) is taken by tens of millions of people worldwide for cardiovascular prevention, stroke prevention, and — in some patients — colorectal cancer prevention. The cardiovascular benefit is well established. The colorectal cancer benefit, demonstrated in several large cohorts, is more interesting because the mechanism has never been entirely clear. Some of it is anti-inflammatory; some of it is anti-platelet; and a growing body of work suggests that some of it may be microbiome-mediated. Aspirin shifts the gut bacterial community, and some of the shifts it produces (toward butyrate-producing taxa, toward a less pro-inflammatory mucosal environment) are in the direction one would expect of a colorectal-cancer-protective intervention.

At the same time, low-dose aspirin causes the same small-bowel injury as the higher-dose NSAIDs, just less of it. The standard combination of low-dose aspirin plus a PPI — protection of the stomach with simultaneous unhindered exposure of the small intestine — may, in fact, *worsen* small-bowel damage by altering the upper gut bacterial environment in a way that exacerbates lower-gut injury. This is one of the paradoxes the field is still working through, and it is one reason why the simple slogan “take a PPI to protect against NSAID damage” is more complicated than it sounds.

What it means in practice

What I want the reader to take away from this section is not that NSAIDs are dangerous and should not be used. They are useful drugs, and for many conditions there is no equally good alternative. What I want the reader to take away is that *long-term* use is the case where the microbiome story matters, and the principle is the same as for PPIs: the lowest effective dose for the shortest necessary duration. An ibuprofen for a sprained ankle is fine. A naproxen taken twice a day for fifteen years for chronic back pain, with no break and no discussion of whether the indication is still strong, is the case where the slow erosion of the small bowel and its bacterial community has had decades to accumulate, and where a conversation with a doctor about whether a different approach might serve the same purpose is worth having. As always, that is a conversation between a patient and a clinician who knows the case — not a recommendation from a popular-science book.

A practical postscript. Several recent animal studies have shown that dietary measures that strengthen the gut barrier — soluble fibre, polyphenol-rich foods, certain fermented foods — reduce NSAID-induced small bowel damage in laboratory models. The human evidence for this protective effect of diet against NSAID enteropathy is, I should say honestly, still thin. But the dietary measures themselves are good for many other reasons, well discussed in Chapter 20, and there is no harm in adding them. If you take an NSAID daily and you also eat thirty different plants a week, you are very probably better off than if you take an NSAID daily and live on white bread and ham. The mechanism by which this is true is not yet pinned down. The direction of the effect, in every animal experiment that has tested it, is the same.

22.6 Statins, Cholesterol, and the Bact2 Enterotype

The most successful drug class of the late twentieth century

Statins are, by any reasonable measure, one of the great therapeutic successes of modern medicine. Atorvastatin, simvastatin, rosuvastatin, pravastatin, and the other members of the class have lowered the cholesterol levels — and, more importantly, the rates of heart attack and stroke — of hundreds of millions of people worldwide over the last three decades. The trial evidence is overwhelming and the cardiovascular benefit in patients at meaningful risk is one of the most reproducible results in clinical medicine.

What is less settled than most people assume is exactly *why* statins work as well as they do. The textbook mechanism is straightforward: they inhibit the enzyme HMG-CoA reductase, the rate-limiting step of cholesterol synthesis in the liver, and the resulting fall in LDL cholesterol slows the progression of atherosclerosis. But if cholesterol reduction were the whole story, then other drugs that lower LDL by a comparable degree ought to produce comparable survival benefits. And here the picture gets interesting. Ezetimibe, which blocks cholesterol absorption from the gut by a completely different pathway, does lower LDL effectively — yet its monotherapy evidence for reducing hard cardiovascular endpoints or overall mortality is remarkably thin; no major trial has even tested ezetimibe monotherapy against placebo for survival [REF:cochrane2019eze]. PCSK9 inhibitors, which drive LDL to impressively low levels by yet another route, do reduce cardiovascular events in large trials, but even here it has taken very large studies to demonstrate the benefit, and all-cause mortality reductions have been elusive [REF:khan2023pcsk9]. The simplest interpretation is that statins derive their clinical benefit from *both* LDL lowering and a suite of cholesterol-independent effects — sometimes called *pleiotropic effects* — that include repair and protection of the vascular endothelium (partly through upregulation of nitric oxide synthase), direct anti-inflammatory activity (statins lower C-reactive protein and suppress several inflammatory cytokines through inhibition of isoprenoid signalling), stabilisation of atherosclerotic plaques, and modest antithrombotic properties [REF:oesterle2017]. The relative contribution of cholesterol lowering versus pleiotropy is still debated, but the debate itself matters for this chapter, because at least some of those pleiotropic effects may turn out to involve the gut microbiome.

What is also true, and rarely discussed outside the specialist literature, is that not everyone responds the same way. A substantial minority of statin users show smaller-than-expected reductions in LDL cholesterol on a given dose, or experience side effects (muscle aches, modest rises in blood sugar) that are sometimes attributed to genetic variation in the enzymes that metabolise the drug. Some of that variability is undoubtedly genetic. Some of it, the more recent literature suggests, is bacterial.

The Bact2 finding

In 2020 a Belgian-led group around Sara Vieira-Silva and Jeroen Raes published a paper in *Nature* that examined the relationship between obesity, inflammation, and the gut microbiome in two large European cohorts [REF:vieirasilva2020]. They had previously characterised what they called *enterotypes* — broad statistical clusters of gut microbial community structure that recur across many human populations. One of those enterotypes, which they named *Bact2*, is characterised by high *Bacteroides* abundance, low microbial cell density, low diversity, and a strong association with chronic inflammation and with several inflammatory and metabolic diseases. Bact2 is, in essence, the “unhealthy” enterotype: it shows up disproportionately in obesity, inflammatory bowel disease, and various metabolic disorders.

The finding that made the paper unexpectedly interesting was this. Among obese individuals, those who were taking statins had a much *lower* prevalence of the Bact2 enterotype than statin-naïve obese individuals with otherwise similar characteristics. The effect was substantial: obese people on statins looked, at least by this microbiome metric, more like lean controls than like other obese people. Statin use, in other words, was associated with a microbiome that was less inflamed.

The paper was careful to state that this was an observational association and not proof of causation — perhaps people with already-healthier guts were more likely to be put on, and tolerate, a statin. But the finding generated a substantial follow-up literature, and the current best interpretation is that statins do exert a real, modest, beneficial effect on the gut bacterial community — partly by changing bile acid composition (statins reduce hepatic cholesterol output, which reduces bile acid production, which changes the microbiome, which then changes how the host handles bile acids in a feedback loop), and partly through direct effects on bacterial growth that have not been fully worked out.

A more complicated picture

More recent work has complicated this story without overturning it, and it leads into a question that I think is genuinely interesting, even if the answer is not yet settled.

A 2022 multi-cohort study by Verhaar and colleagues asked why different people respond so differently to the same statin at the same dose [REF:verhaar2022]. The answer, it turned out, was partly in their gut bacteria. People with a *Bacteroides*-dominated, low-diversity microbiome — essentially the same Bact2 enterotype that Vieira-Silva had linked to inflammation — showed the *strongest* cholesterol-lowering response to statins. They were two to four times more likely to reach their LDL target than people with a more diverse, *Ruminococcus*-enriched gut community. The drug worked hardest in the people whose microbiome was already least healthy.

But here is the twist. Those same Bact2-enterotype individuals — the “best responders” in terms of LDL lowering — also experienced the greatest disruption to glucose control. And the reason appears to be precisely the microbiome: the Bact2 enterotype is depleted in short-chain-fatty-acid producers such as *Faecalibacterium* and *Subdoligranulum*, the bacteria whose metabolic products help maintain insulin sensitivity and nourish the gut lining. People whose gut community was rich in butyrate producers and in *Akkermansia* (the *Ruminococcus*-enriched enterotype) showed statin efficacy *without* the same metabolic cost. The microbiome, in short, was acting as a buffer — or the lack of it was leaving certain patients exposed.

This connects to a broader question in cardiovascular medicine that deserves a mention in this chapter, even though it is not yet a microbiome story with a proven ending. In secondary prevention — that is, in people who have already had a heart attack or stroke — the evidence is clear: more intensive statin therapy saves lives, and driving LDL lower is generally better. But in primary prevention — healthy people being treated to *prevent* a first event — the picture is less straightforward. A 2024 retrospective study of nearly 178,000 primary prevention adults (not on statins, not diabetic) found a striking U-shaped curve between LDL cholesterol and all-cause mortality [REF:lind2024]. The lowest mortality risk sat in the LDL range of 100 to 189 mg/dL. Below that — LDL under 80 mg/dL — mortality was 23 percent *higher* than the reference group, and the association persisted after adjusting for confounders. Very high LDL carried a similar penalty in the other direction [REF:liu2021nhanes]. This is observational data, and confounding by indication (sicker people may have lower cholesterol for reasons that also kill them) is always a concern. But the signal has appeared in enough independent cohorts, including the ASPREE trial of healthy older adults, that

it cannot be easily dismissed.

Nobody has yet proven that the microbiome is part of the explanation for excess mortality at very low LDL levels. But the pieces are suggestive. Statins shift the gut community toward *Bacteroides* dominance and away from SCFA-producing *Firmicutes* [REF:liu2025statin]. The degree of that shift tracks with the magnitude of LDL lowering. The people who respond most aggressively to statins in terms of cholesterol are the ones whose SCFA producers are most depleted. And SCFA depletion is itself associated with increased gut permeability, chronic low-grade inflammation, and impaired glucose metabolism — all of which are risk factors for the very conditions (infection, metabolic disease, frailty) that drive excess mortality at the lower end of the cholesterol distribution. Whether these threads connect causally is a question for the next generation of prospective studies. But I would not be surprised if, in ten years, the U-shaped mortality curve and the statin–microbiome interaction turn out to be part of the same story.

What I take away from the statin literature

Statins are excellent drugs and the cardiovascular benefit, in people who have established cardiovascular disease or who are at genuinely high risk, dwarfs any plausible microbiome cost. That arithmetic is not in doubt. What is in doubt — and what a thoughtful reader should carry away from this section — is whether the same arithmetic applies to the increasingly common practice of prescribing statins for primary prevention in people at modest absolute risk. The microbiome cost is real, modest, and probably partially offsettable by the same kinds of dietary measures (high fibre intake, plant variety) that we recommend independently for cardiovascular health. But it is a cost, and like all costs it matters more when the benefit it is being weighed against is smaller. The statin–microbiome interaction is also a beautiful example of how a drug can be cardioprotective and gut-perturbing at the same time, and of how the simplistic question “is this drug good for you or bad for you?” has to be replaced by the more honest question “what is it doing, and to what, and how much?” If you are on a statin and wondering whether it is right for you, the answer is almost certainly a conversation with your doctor, not a paragraph in a popular-science book — but it is a conversation worth having, especially if the original indication was primary prevention and has never been reviewed.

22.7 The Mind-Altering Drugs: Antipsychotics, Antidepressants, and the Microbial Cost of Calm

Why this section is harder than the others

Of all the drug classes in this chapter, the psychiatric medications — antipsychotics and antidepressants in particular — are the ones I find most difficult to write about responsibly, for two reasons.

The first is that they are taken by many millions of people for serious illnesses that, untreated, cause enormous suffering and substantially increased mortality. Severe schizophrenia, bipolar disorder, major depression, and the more severe forms of anxiety are not minor inconveniences. The drugs that treat them are imperfect and have real side effects, but the alternative — for the patients who actually need them — is often terrible. I do not want a single reader to come away from this section thinking that the gut microbiome is a reason to stop their psychiatric medication. It is not.

The second reason is that the gut microbiome of people with severe mental illness is itself altered, in ways that reflect the illness as well as the treatment. Schizophrenia is associated with a different gut

community than the absence of schizophrenia, even before any drugs are taken. So is depression. So is bipolar disorder. The arrow between gut and brain runs in both directions, as we saw at length in Chapter 13 on the gut–brain axis, and disentangling what the *drug* does from what the *condition* does is genuinely difficult. Most of the relevant studies cannot do it cleanly.

With those caveats acknowledged, let me describe what is known.

Antipsychotics and the weight problem

The atypical antipsychotics — olanzapine, clozapine, risperidone, quetiapine, and several newer agents — have been a substantial improvement over the older drugs in many respects. They have, however, a notorious side effect: weight gain. Some patients on olanzapine gain ten or fifteen kilograms within the first year of treatment. The weight gain is associated with insulin resistance, raised blood lipids, and a substantially increased long-term risk of type 2 diabetes and cardiovascular disease. For many years the standard explanation involved the drugs' effects on appetite-regulating receptors in the brain, particularly histamine and serotonin pathways. This explanation was not wrong, but it was incomplete.

Around 2015 several groups began to notice that olanzapine produces large and reproducible changes in the gut microbiota of laboratory animals, and that these changes track closely with the metabolic effects. A 2024 paper from Kamath and colleagues compared two atypical antipsychotics, olanzapine and lurasidone, in rats over twenty-one days [REF:kamath2024]. Olanzapine, as expected, produced substantial weight gain, raised inflammatory markers, raised blood glucose, and raised triglycerides — and also produced a marked dysbiosis with reduced diversity, depleted short-chain-fatty-acid-producing bacteria, and reduced *Akkermansia muciniphila*. Lurasidone, in the same experimental setup, did almost none of these things. The two drugs had very similar therapeutic profiles in the brain — both occupy dopamine D2 receptors in roughly the same way — but only one of them was wrecking the gut.

This is a useful piece of evidence because it begins to separate the microbiome story from the brain-receptor story. Whatever the antipsychotic is doing in the gut is not a necessary consequence of antipsychotic action; it is a consequence of olanzapine specifically, mediated by mechanisms that lurasidone, despite its similar therapeutic profile, does not share. A 2023 longitudinal rat study tracked the trajectory of olanzapine-induced dysbiosis dose by dose and showed that the abundance of certain *Alistipes* and *Oscillibacter* taxa correlated, almost in lock-step, with the magnitude of weight gain [REF:qian2023olanzapine]. The relationship was tight enough that the gut community was effectively a real-time biomarker of the metabolic damage being done.

The mechanistic question that follows is whether the dysbiosis *causes* the weight gain or is merely a by-product of it. Several papers have addressed this. The most provocative finding is that prebiotic interventions that restore *Akkermansia muciniphila* — including a galacto-oligosaccharide supplement called B-GOS that selectively feeds the species — reduce or prevent the metabolic side effects of olanzapine in mouse models [REF:zeng2024olanzapine]. If the dysbiosis were merely a passive marker, fixing it should not help. Fixing it, in animals, helps. The interpretation, increasingly accepted in the field, is that olanzapine causes some of its weight gain by directly suppressing protective gut bacteria (especially *Akkermansia*), which loses the host the metabolic benefits those bacteria normally provide, and that restoring those bacteria can substantially blunt the weight gain.

A 2024 mechanistic review from Bertossi laid out this case in detail and proposed *Akkermansia* probiotics as an adjunct to olanzapine therapy [REF:bertossi2024]. A 2024 systematic review of

the broader literature, looking at twelve clinical and preclinical studies of various antipsychotics, confirmed that dysbiosis is a near-universal finding across the class but that the magnitude varies enormously from drug to drug [REF:dias2024]. The clinical implication, still being worked out in human trials, is that the choice between antipsychotics — when several agents would do roughly the same therapeutic job — might reasonably take into account which one is likely to do less damage to the gut. This is a conversation between psychiatrist and patient. It is not a reason to refuse a drug that the patient genuinely needs. But for some patients, for whom several options exist, it is a real consideration that did not exist a decade ago.

Antidepressants and the dysbiosis question

The antidepressants are a more complicated case, partly because the class is more heterogeneous and partly because depression itself is associated with gut microbiome changes that make the treatment-versus-illness question hard to answer.

Several SSRIs (selective serotonin reuptake inhibitors) — sertraline and fluoxetine in particular — were among the drugs flagged in the original Maier 2018 screen as having direct antibacterial activity at concentrations comparable to those reached in the gut during normal therapeutic use [REF:maier2018]. This was a surprising finding, because nobody had designed these drugs as antibiotics, and yet there was the antibacterial effect on the dish. A handful of subsequent papers have shown that long-term SSRI use does shift the gut microbiome in human cohorts, generally in the direction of reduced diversity and altered tryptophan metabolism. (Tryptophan, the amino acid precursor of serotonin, is also the substrate from which gut bacteria produce a range of important signalling molecules, so anything that touches tryptophan availability is likely to ripple through both the brain and the gut community.)

What is much harder to say, with current evidence, is whether the microbiome changes induced by SSRIs contribute to the antidepressant effect, oppose it, or are merely incidental. There is a coherent and increasingly evidenced argument — drawing on the gut–brain axis literature in general — that the gut bacteria participate in mood regulation through several routes, and that drugs which affect both the central nervous system and the gut microbiota are likely to be operating on both fronts simultaneously. The metformin-as-antidepressant chain I described earlier in this chapter [REF:xie2024metformin] is a striking example of a drug working on the brain through the gut, and there is no fundamental reason to assume that an SSRI does not do something similar, in addition to its better-known direct effects on serotonin reuptake.

What I would say to a reader on a long-term antidepressant is exactly what I would say to a reader on a long-term PPI: the drug is doing something to your gut community, the something is real and probably modest, and the right question is not “should I stop?” but “is this drug still serving the purpose it was started for, and is the dose still the lowest one that achieves the goal?”. For someone whose depression is being well controlled by an SSRI, the gut microbiome story is a piece of context, not a reason for action. For someone who has been on an SSRI for fifteen years without the indication being formally re-reviewed, that conversation is probably worth having — and it is the same conversation people are increasingly being encouraged to have anyway, on grounds entirely separate from the microbiome.

A careful word about the gut–brain feedback loop

There is one more thing worth saying here, because it is genuinely interesting and because it complicates the picture in a constructive direction. If the gut microbiome influences mood — and the

evidence from Chapter 13 is that it does — and if certain dietary interventions and certain probiotics can shift the gut community in ways that improve mood (the so-called *psychobiotic* literature, which is small but real) — then there is a possibility, still being explored, that for *mild* mood disturbance some of the work that an antidepressant is currently doing pharmacologically might be doable nutritionally instead. I want to be very careful about this. It is not a recommendation. The evidence for psychobiotic intervention is currently nowhere near the evidence for SSRIs in moderate-to-severe depression, and substituting one for the other in a serious depression would be irresponsible and probably dangerous. But for mild low mood, mild anxiety, the kind of thing that exists on a spectrum and not as a clinical diagnosis, the case for trying dietary and lifestyle modifications first — fibre, fermented foods, plant variety, exercise, sleep, the other levers from Chapter 21 — is reasonably strong on its own merits, and the microbiome story adds a small additional reason to do so.

For anything more serious than that, please talk to your doctor, and please do not stop or alter psychiatric medication on the basis of this chapter or any other chapter of this book.

22.8 Opioids and the Bowel That Forgets to Move

The opioid drugs — morphine, codeine, oxycodone, hydromorphone, fentanyl, tramadol, buprenorphine, methadone, and many others — are powerful analgesics and are also, by any reasonable measure, the medication class with the most complicated public health story of the early twenty-first century. They are essential for the management of severe acute and cancer pain. They are problematic for the management of chronic non-cancer pain, where the long-term harm-benefit ratio is much less favourable than was once believed. They are dependent-forming, they suppress respiration, and the period from roughly 1995 to 2020 saw, in several countries, an opioid prescribing pattern whose downstream consequences will be the subject of public health discussion for decades.

I am going to set the prescribing-policy debate aside and talk only about what opioids do to the gut microbiome, because that is the topic of this chapter.

The first thing to say is that nearly everyone who takes a regular opioid develops some degree of *opioid-induced constipation* — OIC for short — and this is recognised by every clinician who treats pain. The proximate mechanism is straightforward and has been understood for decades: opioids bind to mu-opioid receptors on the enteric nervous system (the network of neurons embedded in the wall of the bowel) and reduce the release of acetylcholine, which is the principal driver of gut motility. The bowel slows down. Stool sits in the colon longer. Water is reabsorbed. The result is the hard, infrequent stools that almost every chronic opioid user has to manage.

What is much less appreciated is that the slowed transit, combined with direct effects of opioid receptor activation on the gut epithelium, also produces a substantial dysbiosis. The published human cohort data are sparse — far sparser than for PPIs or metformin — but the animal evidence is consistent and the mechanism is plausible. Chronic morphine administration in mice reduces gut microbial diversity, depletes short-chain-fatty-acid-producing bacteria, weakens the intestinal barrier, and increases bacterial translocation from the gut lumen into the lamina propria and from there into the lymphatic and portal circulation. Some of the systemic side effects of long-term opioid use — including subtle immune dysregulation, increased susceptibility to certain infections, and chronic low-grade inflammation — are now believed to be at least partly mediated by this gut barrier

dysfunction.

The clinical implication is mostly that opioid-induced constipation is both more difficult to manage and more biologically consequential than was once thought. The standard interventions — increased fluids, dietary fibre, osmotic laxatives, and the newer peripherally-restricted opioid antagonists such as naloxegol and methylnaltrexone, which block opioid receptors in the gut without affecting analgesia in the brain — all have a place. What the microbiome literature adds is a reminder that simply forcing the bowel to move is not the whole story; the underlying community has been disrupted, and dietary support of that community (the by now familiar fibre-and-variety formula) is reasonable adjunctive care.

I want to stress that the opioid-microbiome literature in humans is genuinely thin. Almost everything I have just said is supported by mechanistic plausibility and animal experiments. The human cohort studies that would settle the question — long-term opioid users versus matched controls, sequenced longitudinally, with proper attention to the confounders of pain itself, of the conditions causing the pain, of co-prescribed laxatives, and of diet — have not been done at the scale that has been done for metformin or PPIs. I would expect this to change in the next several years. For now, the honest summary is that opioids do appear to cause meaningful gut dysbiosis, the mechanism is well-supported in animals, and the human data are catching up.

22.9 The Forgotten Effect of Laxatives

Of all the entries in this chapter, this is the one I almost left out. Not because laxatives are unimportant — they are taken regularly by a substantial fraction of older adults, and chronically by some — but because the published literature on laxative effects on the gut microbiome is, frankly, surprisingly small. There is one really good paper, and not very many others.

The really good paper is from Carolina Tropini and colleagues working with Justin Sonnenburg at Stanford, published in *Cell* in 2018 [REF:tropini2018]. Tropini's group asked an interesting question: what happens to the structure of the gut microbial community when an osmotic laxative — polyethylene glycol, the active ingredient in many bowel-prep solutions and in chronic-constipation formulations — is administered to mice? The work is best known among microbiome researchers for the elegance with which it disentangled what an ecological catastrophe does to a microbial community. Polyethylene glycol works by drawing water osmotically into the bowel lumen; this dilutes the contents, shortens transit time, and mechanically washes a substantial fraction of the resident bacteria out of the system.

What Tropini and colleagues showed is that this is not a small or transient effect. PEG administration produced a marked drop in microbial cell density and a substantial change in community composition that did not fully recover even weeks after the laxative was withdrawn. Some bacterial species rebounded; some did not. The community that re-established itself after the wash-out was not quite the same as the community that had been there before. In an ecological sense, the laxative had reset the system in a way that had lasting consequences for which species came to dominate the recovered community.

This was a single, careful, mouse-based paper. The directly comparable human work has not, to my knowledge, been done at the same level of detail. But the principle that Tropini's paper established — that even an apparently mild perturbation like an osmotic laxative can produce lasting community-level changes by transiently changing the physical environment — is a useful reminder

that the gut community is shaped not only by what is in it but by the conditions under which it lives. Speed up the transit dramatically and you change which species can hold on. Slow it down dramatically — as opioids do — and you change which species can grow.

A more recent and slightly different angle is the prebiotic effect of *lactulose*, a non-absorbable sugar widely used as both a laxative and (at lower doses) as a treatment for hepatic encephalopathy. Lactulose is fermented by gut bacteria (it is, after all, a sugar) and selectively expands populations of *Bifidobacterium*, in a manner that overlaps substantially with the prebiotics discussed in Chapter 20. So lactulose, used long-term, is doing two things at once: it is moving the bowel, and it is feeding a particular subset of the bacterial community. For some patients — particularly those with hepatic encephalopathy — this dual action is part of the therapeutic intent.

What about chronic stimulant laxatives — senna, bisacodyl, the herbal preparations sold for “regularity”? The published microbiome data on long-term stimulant laxative use are, I have to admit, almost nonexistent. The mechanistic concern is reasonable — chronic alteration of transit time will, like chronic alteration of any major physical parameter, eventually shift which microbes can do well in the bowel — but I cannot point to a definitive paper that has measured this in humans. If you are taking a stimulant laxative every day for years, the same general principle that runs through this chapter applies: minimum effective dose, minimum necessary duration, periodic review with the doctor, and the diet-and-fibre adjuncts that are good for the gut anyway. Chronic constipation has many causes and several treatments; for many patients, addressing the underlying cause (insufficient fluid, insufficient fibre, low physical activity, certain medications including some of the ones earlier in this chapter) is a better long-term strategy than indefinite laxative use. That, again, is a clinical conversation, not advice from a book.

22.10 Hormones in a Pill: Contraceptives, HRT, and the Estrobolome

What the gut microbes do to your hormones

There is a curious feature of human biology that does not get the attention it deserves outside specialist circles. The gut bacteria participate, directly and continuously, in the hormonal economy of the body — particularly in the metabolism of oestrogens. This is the basis of a concept that the microbiologist Claudia Plottel and her colleague Martin Blaser introduced in 2011 and called the *estrobolome*: the collection of bacterial genes in the gut whose products metabolise oestrogens.

The mechanism, in outline, is this. Oestrogens produced by the ovaries (or, after the menopause, by the adrenal glands and adipose tissue) are conjugated in the liver — that is, they have a sugar group called glucuronide attached to them, which inactivates them and earmarks them for excretion through the bile. The conjugated oestrogens travel down the bile duct into the duodenum and from there into the small bowel, where, in a healthy adult, certain gut bacteria express an enzyme called β -glucuronidase that cleaves the glucuronide group off the oestrogen molecule and reactivates it. The reactivated oestrogen is then reabsorbed across the gut wall and re-enters the circulation, where it does its hormonal work all over again. The total amount of oestrogen the body experiences over the course of a day is therefore not just the amount the ovaries (or adrenals) made; it is that amount, plus everything the gut bacteria sent back into circulation. The estrobolome is, in effect, a hormonal thermostat with a microbial component.

This becomes interesting for the present chapter because oral contraceptives and hormone replacement therapy both deliver pharmacological oestrogen (and, in most contraceptives, a progestogen

as well) into the system, and that pharmacological hormone interacts with the same enzymatic machinery as the body's own. A daily oestrogen pill is, from the point of view of the estrobolome, an additional hormonal substrate for the bacteria to handle. And — perhaps unsurprisingly — the gut microbiome of women on long-term oral contraceptives differs measurably from that of women not on them.

The published evidence

The strongest recent human work in this area is a small but careful longitudinal study by Hua and colleagues, published in 2022 in the *Journal of Medical Microbiology* [REF:hua2022]. They followed ten women across the initiation of oral contraceptive therapy, sequencing the gut microbiome before, one month after, and six months after starting the pill. They did not find dramatic shifts in overall diversity or in the dominant taxa, which was somewhat reassuring. What they did find was a measurable reorganisation of the microbial functional capacity: increases in the bacterial pathways for synthesising peptidoglycan (the cell wall component of Gram-positive bacteria), several amino acids, and the salvage pathway for nicotinamide adenine dinucleotide. Serum oestradiol and SHBG (sex hormone binding globulin, the carrier protein that controls free oestrogen levels) were associated with the abundance of *Bacteroides* species capable of degrading flavonoids — plant-derived compounds with weak oestrogenic activity. The picture, in short, was of a gut community that was *functionally* responding to a new hormonal input even when its taxonomic composition looked relatively unchanged.

Subsequent work in 2025 from several groups has added to the picture. Studies in physically active women on hormonal birth control have reported reductions in beta diversity and in short-chain-fatty-acid-producing bacteria compared with non-users [REF:brito2025]. A separate study in mice on longer-term oral contraceptive exposure showed segment-specific changes (the colon was the most affected part of the bowel) that accumulated over time [REF:clapp2025]. The general direction is that hormonal contraceptive use is associated with modest but real shifts in the gut community, that the shifts accumulate over months and years rather than appearing within days, and that they preferentially affect short-chain-fatty-acid producers — the same group of bacteria that turn up in almost every entry of this chapter as a marker of a healthy gut.

What this means for users

I want to be quite clear about what this section is and is not saying.

It is not saying that oral contraceptives cause harm to the gut microbiome serious enough to outweigh their benefits as contraception. The benefits of effective contraception — reduced unintended pregnancy, reduced maternal mortality, reduced reliance on more invasive interventions — are very large. The microbiome cost is real but, on current evidence, modest. For the great majority of users, the trade-off comes out clearly in favour of using the drug.

It is also not saying that hormone replacement therapy at the menopause is something to avoid. HRT, for women whose menopausal symptoms are sufficiently disruptive, is a reasonable and often excellent treatment. The microbiome question for HRT is even less well studied than for contraceptives, but the mechanism is the same — oestrogens cycling through the gut and through the estrobolome — and the magnitude of the effect appears to be modest in the same direction.

What this section is saying is that for women on long-term hormonal contraception or HRT, the standard advice that runs through this whole chapter applies: keep dietary fibre intake high, eat

a varied range of plants, include fermented foods if you tolerate them, and have a periodic conversation with your prescribing doctor about whether the dose and the duration are still appropriate for your situation. The microbiome is one of several considerations on the larger ledger of these medications, and like the other considerations, it is best discussed openly rather than ignored.

There is one further interesting wrinkle worth mentioning. Because the gut bacteria recycle oestrogens, the *composition* of a woman's gut microbiome may influence how much oestrogen her body actually experiences from a given dose of an oral contraceptive. This is, at the moment, a research question rather than a clinical recommendation — the variation between women is substantial, and the clinical relevance has not been pinned down — but it is part of the broader emerging story that hormonal pharmacology and the gut bacteria are talking to each other in ways that the original drug developers did not anticipate. It is also one more illustration of the chapter's recurring theme: the same drug, in different people, may be doing different things, partly because of who lives in their gut.

22.11 The Supplement Aisle: Iron, Vitamin D, Fish Oil, and What's Real About Berberine

A note on “natural”

Before I begin this section, a quick word on the framing. Many readers will, I suspect, take supplements regularly — perhaps a multivitamin, perhaps fish oil, perhaps vitamin D, perhaps a herbal preparation recommended by a friend or read about online. The cultural assumption around supplements is that because they are “natural” or “from food”, they are categorically different from drugs and could not possibly have the kinds of effects on the body that pharmaceuticals have. This is not quite right. A supplement that is biologically active enough to do anything at all is, by definition, biologically active. The molecules it contains do not know whether they came from a pharmaceutical company or a herbalist's shelf. The gut microbiome, in particular, treats the two with complete impartiality. A daily oral iron tablet is, from the point of view of the bacteria in your duodenum, an iron tablet — they do not care whether the box it came in had a prescription label or a “wellness” label.

With that established, let me look at the four supplement categories the evidence has the most to say about: iron, vitamin D, omega-3 fish oil, and the one herbal medicine that has accumulated genuinely interesting data, berberine.

Oral iron and the *Enterobacteriaceae* bloom

Iron deficiency is one of the most common nutritional problems in the world, particularly among menstruating women, pregnant women, infants, and older adults. The standard treatment — a daily oral iron tablet, typically ferrous sulphate or ferrous fumarate, often in doses of 60 to 200 mg of elemental iron — is cheap, generally effective, and prescribed (or self-prescribed) in enormous quantities. What the published microbiome literature has shown over the past decade, with reasonable consistency, is that an oral iron supplement of this kind has a measurable and not entirely benign effect on the gut bacterial community.

The basic problem is that iron is a contested resource in the gut. Many of the most desirable bacteria in the human colon — the *Bifidobacterium* and *Lactobacillus* species that produce short-chain

fatty acids, that support the gut barrier, and that out-compete pathogens — have relatively low iron requirements. Many of the less desirable bacteria — *Escherichia coli* and other members of the *Enterobacteriaceae* family, which include several known and several conditional pathogens — are, by contrast, voracious iron consumers and have evolved sophisticated machinery (siderophores, specialised transporters) to scavenge iron from their environment. A high-dose oral iron tablet that has not been absorbed in the duodenum (which most of it has not, because oral iron absorption is inefficient) reaches the colon and provides exactly the resource the more pathogenic bacteria are best at exploiting. The result, in several studies in African infants given iron-fortified weaning foods, was a measurable bloom of *Enterobacteriaceae* and a measurable increase in markers of intestinal inflammation. In some studies the iron-supplemented infants also had higher rates of clinical diarrhoea, although this finding has been inconsistent.

A 2024 randomised controlled trial in Bangladeshi infants, the largest of its kind to date, gave a more nuanced picture [REF:baldi2024]. With nearly a thousand children randomised to placebo, oral iron, or iron-containing micronutrient powders, the headline findings on diarrhoea were reassuring — no overall increase in diarrhoea was seen — and the primary microbiome analysis, after adjustment for multiple comparisons, did not show statistically significant changes in alpha or beta diversity. But the unadjusted trends were consistent with the earlier literature: a decrease in *Bifidobacterium* and *Lactobacillus*, and an increase in *Clostridium* species with pathogenic potential, particularly in children whose iron stores were already adequate at baseline. The lesson the field has gradually drawn from this body of work is that oral iron supplementation is most useful and least risky in children (and adults) who are genuinely iron-deficient, and that giving it to people who are not deficient is the case in which the gut microbiome cost looks most concerning.

For an adult reader who is taking an iron tablet, the practical implications are these. First, take iron because you have a documented need for it — a low ferritin, a low haemoglobin with appropriate workup, a clear medical reason — rather than as a routine preventive supplement. Second, the lower-dose, slower-release formulations (such as ferrous bisglycinate, sucrosomial iron, and the alternate-day dosing strategies that have become more popular in the last few years) appear to deliver useful amounts of absorbed iron with substantially less unabsorbed iron reaching the colon, and on first principles should produce smaller microbiome perturbations. Third, iron given by injection — when oral iron is not tolerated or absorbed — bypasses the gut entirely and avoids the colonic-bloom problem. Animal data support this route as gentler on the bowel ecology [REF:tian2025iron]. Most clinicians do not switch to intravenous iron lightly because it is more invasive and more expensive, but it is a reasonable option for patients who have shown they cannot tolerate oral iron well.

Vitamin D: probably good for the gut, although the data are still settling

Vitamin D is a slightly unusual case. Unlike most of the entries in this chapter, the published evidence on vitamin D and the gut microbiome is, if anything, modestly *positive*. Vitamin D deficiency is associated with various forms of dysbiosis in human and animal studies; correcting the deficiency tends to improve the gut community and to strengthen the gut barrier. A 2025 mouse study showed that withdrawing vitamin D in an aging-model system worsened intestinal permeability, depleted *Akkermansia muciniphila*, and shifted the community in a pro-inflammatory direction [REF:feng2025vitd]. Another 2025 paper showed that vitamin D supplementation improved the gut community and the hypothalamic-pituitary-adrenal axis in mice subjected to sleep desynchrony [REF:wu2025vitd]. The mechanism appears to involve direct effects of vitamin D on epithelial barrier proteins, on intestinal antimicrobial peptide production, and on the regulation of various aryl

hydrocarbon receptor (AhR) signalling pathways that the gut bacteria themselves use.

The clinically actionable point, for a reader who has been told they are vitamin D deficient and offered a supplement, is that the vitamin D is probably doing their gut a small favour as well as the more famous favour to their bones. The evidence base does not yet support megadose supplementation as a microbiome intervention in someone whose vitamin D level is normal — the relevant studies have not been done. The standard message of “correct deficiency, do not chase ever-higher levels” applies here as it does to vitamin D supplementation more generally.

Fish oil and the omega-3 question

The omega-3 fatty acids found in fish oil — eicosapentaenoic acid (EPA) and docosahexaenoic acid (DHA) — have been the subject of an enormous research and supplementation industry over the past three decades. Their cardiovascular benefits in healthy populations are smaller and more controversial than was once claimed, but their anti-inflammatory effects are real, and several animal studies have shown that they shift the gut microbiome in ways that look broadly favourable: increased *Akkermansia muciniphila*, increased *Bifidobacterium*, increased *Lactobacillus*, increased short-chain fatty acid production, and reduced abundance of certain pro-inflammatory taxa. A 2025 systematic review of omega-3 effects on the gut microbiome in animal models of obesity collated thirty-two trials and found a remarkably consistent picture across this whole literature [REF:cao2025omega3].

A particularly elegant 2025 paper using a simulated human intestinal model showed that gut microbes themselves metabolise EPA and DHA into a class of anti-inflammatory lipid mediators called *oxylipins* — molecules with documented anti-inflammatory and anti-cancer properties [REF:roussel2025fishoil]. So the fish oil is not just a passive substrate; it is being converted by the bacteria into bioactive products that then act on the host. This is one of those mechanistic chains that makes the simple question “does fish oil work?” turn into the more interesting question “what does fish oil get turned into, by whom, and what does that do?”.

The human evidence, I should say, lags the animal evidence. Most of the cleanest mechanistic work is in laboratory models; the human cohort and randomised trial data on omega-3 supplementation and the gut microbiome are smaller and less consistent. For the moment my honest summary is that fish oil at modest doses (one to three grams a day of combined EPA and DHA) appears to be a reasonable addition to most diets, has a small favourable effect on the gut microbiome on the available evidence, and is not going to do harm. Whether it produces the cardiovascular benefit it was once heavily marketed for is a separate question on which the evidence has cooled. The microbiome side of the story is one of the better arguments left for it.

Berberine: the herbal medicine that turned out to be real

Most herbal supplements have, when subjected to careful testing, turned out to be either inert, modestly active, or active in ways that depend so heavily on the preparation method as to be useless to recommend. Berberine is the unusual case of a traditional medicine — a yellow alkaloid extracted from the bark of *Berberis* species and used in Chinese and Ayurvedic medicine for centuries — that has stood up to modern pharmacological scrutiny rather well, and that has a particularly interesting microbiome story.

Berberine is poorly absorbed from the gut into the bloodstream — only a few percent of an oral dose reaches the systemic circulation — which historically was considered a problem for its development

as a drug. What several research groups have shown over the past decade is that the poor systemic absorption is in fact part of the reason it works. Most of the oral dose reaches the colon, where it interacts directly with the gut microbiota, suppresses certain bacterial groups, expands others, and produces effects that parallel those of metformin — increased *Akkermansia muciniphila*, increased short-chain fatty acid production, improved bile acid metabolism, improved insulin sensitivity, and modest reductions in fasting blood glucose and lipids [REF:guo2023berberine] [REF:kong2020berberine]. Several head-to-head animal comparisons have shown berberine to be approximately as effective as metformin for some metabolic endpoints, with somewhat better effects on lipids and somewhat weaker effects on glucose.

I should be clear about what berberine is *not*. It is not a tested or licensed treatment for type 2 diabetes in any major regulatory framework, and the human randomised trial evidence, although it exists, is much smaller than for metformin. It is also not free of effects — gastrointestinal upset is common at higher doses, and there are several documented drug interactions, particularly with drugs that are metabolised by the cytochrome P450 system. I would not recommend that any reader replace a prescribed metformin with self-administered berberine on the basis of this paragraph or any other paragraph.

What I would say is that the berberine story is a useful corrective to the assumption that traditional herbal medicines are uniformly ineffective or that microbiome-mediated mechanisms are necessarily weaker than direct pharmacological ones. Berberine works partly because it stays in the gut. Its mechanism is its supposed limitation. It is a reminder that the gut bacteria are a legitimate pharmacological target, that the molecules required to reach that target do not need to be designed by pharmaceutical companies, and that we are very probably going to discover more drugs (or rediscover more old ones) whose action runs through the microbiome rather than around it.

The other commonly-promoted herbal compounds — curcumin, resveratrol, ginseng, the various polyphenol-rich extracts marketed for “gut health” — have fragmentary microbiome data, mostly from animal experiments, with mechanistic stories that are plausible but not yet supported by serious human trials. I would not write any of them off; I would also not pay a lot of money for them on the strength of the existing evidence. The most reliable way to deliver this class of compound to the gut microbiome remains the way humans have always done it: by eating a varied diet of whole foods, particularly plants, particularly the more brightly coloured and bitter ones. That recommendation, made for entirely separate reasons in Chapter 20, turns out to be the sensible default here as well.

22.12 Polypharmacy: The Cumulative Tax

So far, every section of this chapter has discussed a single drug class as if it were the only medication a person was taking. For an unusual reader — a healthy thirty-year-old on a single prescription — that simplification is reasonable. For a much larger group of readers, particularly those over sixty, it is not. The average eighty-year-old in many developed countries is on five or more chronic medications. The average resident of a long-term care facility is on seven or more. Every additional prescription is a separate small intervention on the gut ecosystem, and the question is whether those interventions add up linearly, whether they cancel each other out, or whether they produce something worse than the sum of their parts.

The most serious attempt to answer this question was a 2020 paper by Arnau Vich Vila and colleagues, working with Alexandra Zhernakova and Cisca Wijmenga in the Netherlands, that analysed

the gut microbiomes of nearly two thousand people in the LifeLines DEEP cohort and explicitly examined the effects of multiple concurrent medications [REF:vichvila2020]. The Vich Vila paper was important for several reasons. It confirmed the Maier 2018 finding in a large human cohort: many non-antibiotic drugs do, in fact, perturb the gut community in real life and not only in the dish. It identified specific drug combinations that had synergistic effects on the microbiome — that is, the combination produced more dysbiosis than would have been predicted from adding up the effects of the drugs taken separately. And it established that the cumulative *number* of medications a person was on was, by itself, a meaningful predictor of microbiome disturbance, even after adjusting for the specific drugs involved.

This last finding is the one I want to dwell on, because it has practical implications. A person on eight medications has, on average, a more disturbed gut community than a person on three, even if the drugs are individually mild. The community disturbance is not produced by any one of the drugs being severely toxic; it is produced by the cumulative tax that each modest perturbation extracts. And the same population — older adults on multiple chronic medications — is also the population most vulnerable to the consequences of dysbiosis: increased risk of *Clostridioides difficile* infection, increased risk of bacterial translocation through a leaky gut, altered drug metabolism in unpredictable directions, and a slower recovery from any further perturbation. The biological insult and the biological vulnerability are concentrated in the same group of patients.

A 2024 paper from Garcia-Santamarina and colleagues, published in *Cell*, took a different and equally interesting angle [REF:garcia2024]. Using a synthetic community of thirty-two human gut bacterial species in defined culture, they exposed the community to thirty different drugs and measured the consequences. What they found was that bacterial communities have a fair amount of resilience — drug-sensitive species were often protected, in the community context, by drug-resistant neighbours that biotransformed or detoxified the drug locally before it could reach them. Cross-protection between species was observed about six times more often than cross-sensitisation. The implication is that a healthy, diverse gut community can absorb a substantial amount of pharmacological perturbation without collapsing, because it is not a collection of isolated species but a network of interacting metabolisms. The other implication, more sobering, is that this resilience is finite: when the perturbation is large enough, or when the diversity has already been reduced (as it has in elderly polypharmacy patients, in people recovering from antibiotics, in people with inflammatory bowel disease), the network can no longer compensate, and the community collapses in a way that a single-species view of pharmacology would not have predicted.

The clinical takeaway from the polypharmacy literature is, I think, the most actionable one in this whole chapter. It is the practice of *deprescribing* — the systematic, periodic review of long-term medications to ask whether each one is still indicated, whether the dose can be reduced, and whether some can be safely stopped. Deprescribing is a real and growing area of geriatric medicine, with its own guidelines and a small but rapidly expanding evidence base. Most of the existing arguments for deprescribing are made on grounds entirely separate from the microbiome — falls, cognitive impairment, drug-drug interactions, simple cost. The microbiome story adds one more reason. Each additional drug taken without continuing necessity is, among other things, a small additional tax on the gut community, and the gut community is already under more pressure in the people most likely to be on too many drugs.

If you are in your seventies or eighties and on five or more medications, and you have not had a formal medication review in several years, that review is worth requesting. It is, again, a conversation between you and your doctor, and there is no implication that the answer will be to stop anything. The likelihood is that some of the medications will turn out to be still necessary, some

will turn out to be candidates for dose reduction, and one or two may be relics of older prescribing decisions whose original indication is no longer present. Each one that comes off the list is a small piece of pressure removed from the ecosystem inside you.

22.13 Pharmacomicrobiomics: When Your Microbes Decide What Your Medicine Does

A two-way street

So far in this chapter the story has run in one direction: drugs do things to bacteria. There is, however, an equally important and equally interesting story running in the opposite direction. Bacteria do things to drugs. The gut microbiome is a chemical factory the size of a large pharmaceutical company, and many of the drugs you swallow are processed — sometimes activated, sometimes inactivated, sometimes converted into entirely different molecules — by bacterial enzymes before they ever reach their intended target. The young field that studies this is called *pharmacomicrobiomics*, and although it is still in early days, several of its findings already have direct clinical relevance.

The landmark paper that put the whole question on the map was from Michael Zimmermann, Maria Zimmermann-Kogadeeva, and Andrew Goodman at Yale, published in *Nature* in 2019 [REF:zimmermann2019]. Their team took 271 different oral drugs in current human use and tested each one against a panel of human gut bacterial isolates, asking a simple question: how many of these drugs are chemically modified by gut bacteria? The answer was 176 of the 271 — about two thirds. The bacteria were doing something chemically interesting to the majority of the drugs the researchers tested. Some of those modifications were minor; some were dramatic, and dramatically affected the activity of the drug. The Zimmermann team also identified, in many cases, the specific bacterial enzymes responsible — opening up the possibility of predicting, from a person's gut metagenome alone, how their bacteria are likely to handle a given drug.

The classical examples

A few examples will make the practical relevance concrete.

Digoxin and *Eggerthella lenta*. Digoxin is an old and still useful drug for certain heart conditions, particularly atrial fibrillation and heart failure. It has a notoriously narrow therapeutic window — too little and it does nothing, too much and it is dangerous — and clinicians have long known that digoxin levels in the blood vary unpredictably from patient to patient even at the same dose. Henry Haiser and Peter Turnbaugh, working at Harvard, showed in 2013 that part of this variability is due to a single gut bacterium, *Eggerthella lenta*, which carries an enzyme system that inactivates digoxin in the gut [REF:haiser2013]. People who have abundant *Eggerthella lenta* (and the relevant gene cluster) inactivate a substantial fraction of an oral digoxin dose before it can be absorbed; people who do not, do not. Adding a high-protein meal — which paradoxically appears to suppress *Eggerthella*'s digoxin-inactivating enzyme — increases the bioavailability of digoxin in the affected individuals. This is a bacterium changing how much of a drug actually reaches its target. It is not a rare oddity; *Eggerthella lenta* is a normal member of the human gut microbiota, present in most people, but variable in abundance.

Levodopa and *Enterococcus faecalis*. Levodopa is the cornerstone of pharmacological treatment for Parkinson's disease. It is taken by mouth, absorbed in the small intestine, and converted

in the brain into dopamine, the neurotransmitter that Parkinson's patients are deficient in. A long-standing puzzle has been why the dose required varies so much between patients, and why oral levodopa loses some of its effectiveness over years of treatment. Vayu Maini Rekdal and Emily Balskus, working at Harvard, published a paper in *Science* in 2019 showing that *Enterococcus faecalis* — a normal gut bacterium — carries an enzyme called tyrosine decarboxylase that converts levodopa into dopamine in the gut lumen, before it has a chance to reach the brain [REF:mainirekdal2019]. The dopamine produced peripherally cannot cross the blood–brain barrier, so the drug effectively gets wasted. Patients with high *Enterococcus faecalis* burden need more levodopa to achieve the same brain effect, and that requirement varies between individuals in a way that the specifics of brain chemistry alone do not explain. Maini Rekdal and Balskus also identified specific small-molecule inhibitors of the bacterial enzyme that, in principle, could be co-administered with levodopa to prevent the gut wasting. A clinical implementation of that idea has not yet emerged, but the proof of concept is there.

Irinotecan and bacterial β -glucuronidases. Irinotecan is a chemotherapy drug used for several cancers, including colorectal cancer. After it does its job in tumour cells, the drug is conjugated in the liver and excreted into the bile in an inactive form. In the colon, gut bacteria expressing a particular family of enzymes — bacterial β -glucuronidases — cleave the conjugate and reactivate the drug locally. The reactivated drug then damages the colonic epithelium, causing the severe diarrhoea that is one of the most feared side effects of irinotecan therapy. This is a case where a bacterial enzyme is essentially producing the toxicity of a drug. Several research groups are now developing inhibitors of bacterial β -glucuronidases as adjuncts to irinotecan therapy, with the explicit goal of preventing colonic reactivation while leaving the drug's anti-tumour effect intact. Some of these inhibitors are now in early clinical testing.

The bigger picture

I have given three examples because they are well characterised and clinically vivid. What the field is increasingly suggesting is that these examples are not isolated curiosities; they are the visible peaks of a much larger underwater landscape of bacterial drug metabolism that is only beginning to be mapped. The Zimmermann 2019 paper — bacteria modifying two thirds of 271 tested drugs — implies that the digoxin and levodopa stories are typical rather than exceptional, and that for many drugs the story we currently tell about pharmacokinetics in textbooks (the ADME mnemonic — absorption, distribution, metabolism, excretion — that medical students learn) is missing a crucial fifth letter, M for microbial.

The clinical implication, looking forward, is that the same drug at the same dose can produce very different blood levels and very different therapeutic effects in two patients, depending on what is in their gut. This is the same lesson the metformin section of this chapter delivered, from the other direction: the gut bacteria are not just a passive recipient of pharmacological intervention; they are an active and unpredictable participant. Personalised medicine, as I noted earlier, has so far focused mostly on the human genome. The pharmacomicrobiomics literature is making it harder and harder to ignore the bacterial genome.

For the moment, the practical implications for an individual reader are limited. There is no commercially available “gut metagenome panel” that will tell you which drugs you are likely to handle differently because of your bacteria, and the field is some years away from being able to use such information to adjust prescribing in real time. But the principle is worth knowing about, both because it is genuinely interesting and because it provides a deeper understanding of what is going on

when a drug works wonderfully in your neighbour and not at all in you. The textbook explanation — “individual variation” — has, for at least some drugs, a specific cause, and the cause lives in your gut.

22.14 What This Means for the Person Reading This Book

Let me try to gather the threads of this long chapter into a few practical, undramatic conclusions.

The first conclusion is that almost every drug you can swallow does something to the gut microbiome. For a quarter of common non-antibiotic drugs, that something is direct suppression of bacterial growth in the dish, at human-relevant concentrations [REF:maier2018]. For the rest, the effects are usually subtler — shifts in community composition, changes in metabolite production, alterations in the activity of specific bacterial enzymes, perturbations of the gut barrier — but they are real, they are measurable, and they are not zero. The era in which it was reasonable to assume that a drug acts only on its declared target is over.

The second conclusion is that this fact is, for the most part, *not* a reason to stop your medication. The drugs that have made it through modern regulatory processes did so because, on balance, their benefits outweigh their harms. The microbiome cost is one harm among several, and in the great majority of cases it is one a sensible patient and a sensible doctor will accept in exchange for the benefit on offer. A blood-pressure pill that prevents your stroke is not a less good idea because it has been found to nudge your *Bacteroides* count by ten percent. A statin that prevents your heart attack is not a less good idea because it modestly depletes some butyrate producers. The arithmetic, for the indications for which these drugs are properly prescribed, comes out clearly in favour of taking them.

The third conclusion is that *long-term* use of any chronic medication is the case where the microbiome story matters most, and where the principle of *minimum effective dose for the minimum necessary duration* deserves to be more than a slogan. A short course of any of the drugs in this chapter is a small perturbation that the gut community is well-equipped to recover from. Decades of daily exposure, particularly when the original indication has faded or was never strong, is a sustained pressure that the community handles less well. The actionable response is not to stop the drug; it is to have a periodic conversation with your doctor about whether the indication still holds, whether the dose can be reduced, and whether the continuing benefit is still clearly worth the continuing cost.

The fourth conclusion concerns supplements. A supplement that does anything biologically also does something to the gut bacteria, and the cultural assumption that “natural” equals “harmless” does not survive contact with the data. Iron supplements, in particular, deserve more respect than they often get; they should be taken when there is a documented need and not as a routine precaution. Vitamin D and fish oil look reasonable on the available evidence and are worth considering when the indication is appropriate. The herbal aisle is mostly hype, with berberine as the interesting exception that has done well in head-to-head testing for some metabolic indications. Most of the gut microbiome benefit that supplement marketers promise from their bottles is more reliably delivered by eating a varied diet of plants — a recommendation made for entirely separate reasons in Chapter 20.

The fifth conclusion concerns polypharmacy. If you are over sixty and on five or more chronic medications, and you have not had a formal medication review in some years, please ask for one.

The strongest case for it is not the microbiome — it is falls, cognitive function, drug interactions, cost, and the simple principle that medications accumulate over a lifetime more easily than they get removed. The microbiome story adds one more reason in the same direction. Deprescribing is, in many cases, one of the kindest things a clinician can do for a long-term patient, and it is a conversation more worth having than most patients realise.

The sixth and final conclusion is the one I find most interesting personally. The science of pharmacomicrobiomics — bacteria doing things to drugs as well as drugs doing things to bacteria — is making it clear that the same drug, at the same dose, in two patients with apparently identical clinical situations, can do very different things, partly because their gut communities are not the same. This is not a recommendation. There is no test you can take, no action you can perform on the basis of this knowledge today. But it is a way of understanding why the simple picture of pharmacology you may have been taught — drug goes in, target gets hit, effect follows — is, for a substantial fraction of the drugs in use, an oversimplification of something genuinely more complicated and more interesting. Your microbes are, in a small but real way, deciding what your medicine does.

I want to leave the reader with the same image that closed Chapter 21. The body is an ecosystem under continuous chemical pressure, much of the pressure invented in the last hundred years. The medications discussed in this chapter are some of the most effective inventions in the history of human civilisation; the microbiome cost we have only recently learned to see is a real but, in most cases, modest part of the larger ledger. Knowing that the cost is on the ledger does not change whether to use the drugs. It changes how we use them, and how we talk about them, and how we think about the slow, accumulating relationship between modern pharmacology and the trillions of small partners we have been carrying around inside us all along.

The next two chapters turn, finally, from the question of *what we do to the microbiome* to the question of *what we can do for it*. Chapter 23 takes up the long and complicated story of probiotics — the supplements you can buy in any pharmacy that promise to deliver “good bacteria” to your gut, and the rather more careful science of which of these promises is real. Chapter 24 takes up faecal microbiota transplantation, the most extreme and most dramatic of the microbiome interventions, the one that has saved lives from *Clostridioides difficile* and that is now being tested for a rapidly expanding list of other conditions. Together, those chapters complete the picture this chapter opened: now that we know how easily we can perturb the gut community, what can we do, deliberately, to put it back together?

References for Chapter 22

[REF:maier2018] Maier L, Pruteanu M, Kuhn M, et al. Extensive impact of non-antibiotic drugs on human gut bacteria. *Nature*. 2018;555(7698):623-628. doi:10.1038/nature25979 (PMID: 29555994)

[REF:imhann2016] Imhann F, Bonder MJ, Vich Vila A, et al. Proton pump inhibitors affect the gut microbiome. *Gut*. 2016;65(5):740-748. doi:10.1136/gutjnl-2015-310376 (PMID: 26657899)

[REF:shin2002ppi] Shin JM, Sachs G. Restoration of acid secretion following treatment with proton pump inhibitors. *Gastroenterology*. 2002;123(5):1588-1597. doi:10.1053/gast.2002.36593 (PMID: 12404233)

[REF:zhang2025ppi] Zhang M, et al. Proton pump inhibitor increases intestinal epithelial paracel-

- lular permeability via the p38-MAPK/NF- κ B signaling pathway. *Mol Biol Rep.* 2025;53(1):170. doi:10.1007/s11033-025-11336-y
- [REF:xu2025ppi] Xu J, et al. Unraveling gut microbiome interferences in cancer immunotherapy: a meta-analysis of diverse drug effects. *BMC Cancer.* 2025;25:1776. doi:10.1186/s12885-025-15094-9
- [REF:sharma2025ppi] Sharma S, Rawal P, Kalra S. Gut Guardianship: Recommendations for Rational Use of Proton Pump Inhibitors in Chronic Kidney Disease. *J Pak Med Assoc.* 2025;75(11):1822-1824. doi:10.47391/JPMA.25-94
- [REF:forslund2015] Forslund K, Hildebrand F, Nielsen T, et al. Disentangling type 2 diabetes and metformin treatment signatures in the human gut microbiota. *Nature.* 2015;528(7581):262-266. doi:10.1038/nature15766 (PMID: 26633628)
- [REF:su2025metformin] Su W, et al. The Microbial Regulation Spectrum of Metformin in Patients With Type 2 Diabetes: An Individual-Based Meta-Analysis of 1431 Participants. *J Clin Endocrinol Metab.* 2025;110(8):2383-2403. doi:10.1210/clinem/dgaf238
- [REF:attaye2024] Attaye I, et al. Oral Akkermansia muciniphila L2-7 augments glycemic control in type 2 diabetes. *iScience.* 2024;27(8):110455. doi:10.1016/j.isci.2024.110455
- [REF:xie2024metformin] Xie X, et al. Metformin reprograms tryptophan metabolism via gut microbiome-derived bile acid metabolites to ameliorate depression-like behaviors in mice. *Brain Behav Immun.* 2024;123:442-455. doi:10.1016/j.bbi.2024.09.014
- [REF:rosell2024] Rosell-Díaz M, et al. Metformin-induced changes in the gut microbiome and plasma metabolome are associated with cognition in men. *Metabolism.* 2024;157:155941. doi:10.1016/j.metabol.2024.155941
- [REF:tropini2018] Tropini C, Moss EL, Merrill BD, et al. Transient osmotic perturbation causes long-term alteration to the gut microbiota. *Cell.* 2018;173(7):1742-1754.e17. doi:10.1016/j.cell.2018.05.008 (PMID: 29906449)
- [REF:varela2025] Varela-Mendoza AA, et al. Kombucha Prevents Indomethacin-Induced Enteric Damage in Wistar Rat by Enhancing Epithelial Gut Barrier and Modulating Gut Microbiota. *Food Sci Nutr.* 2025;13(8):e70804. doi:10.1002/fsn3.70804
- [REF:tuerxuntayi2025] Tuerxuntayi A, et al. Paeoniflorin attenuates ASA-induced enteropathy in simulated high-altitude hypoxia in rats by regulating intestinal inflammation, intestinal barrier, and intestinal flora. *Toxicol Appl Pharmacol.* 2025;506:117641. doi:10.1016/j.taap.2025.117641
- [REF:cochrane2019eze] Defined SR, Defined SR. Ezetimibe for the prevention of cardiovascular disease and all-cause mortality events. *Cochrane Database Syst Rev.* 2019;2019(11):CD012502. doi:10.1002/14651858.CD012502.pub2
- [REF:khan2023pcsk9] Khan SU, et al. PCSK-9 Inhibitors and Cardiovascular Outcomes: A Systematic Review With Meta-Analysis. *Cardiovasc Ther.* 2023;2023:4653180. doi:10.1155/2023/4653180
- [REF:oesterle2017] Oesterle A, Laufs U, Liao JK. Pleiotropic Effects of Statins on the Cardiovascular System. *Circ Res.* 2017;120(3):573-585. doi:10.1161/CIRCRESAHA.116.308537 (PMID: 28154098)
- [REF:vieirasilva2020] Vieira-Silva S, Falony G, Belda E, et al. Statin therapy is associated with lower prevalence of gut microbiota dysbiosis. *Nature.* 2020;581(7808):310-315.

doi:10.1038/s41586-020-2269-x (PMID: 32433607)

[REF:verhaar2022] Verhaar BJH, Collard D, Prodan A, et al. Heterogeneity in statin responses explained by variation in the human gut microbiome. *Med.* 2022;3(6):388-405.e6. doi:10.1016/j.medj.2022.05.002

[REF:lind2024] Lind L, et al. Is LDL cholesterol associated with long-term mortality among primary prevention adults? A retrospective cohort study from a large healthcare system. *BMJ Open.* 2024;14(3):e080151. doi:10.1136/bmjopen-2023-080151

[REF:liu2021nhanes] Liu X, et al. Association between low density lipoprotein cholesterol and all-cause mortality: results from the NHANES 1999–2014. *Sci Rep.* 2021;11(1):22111. doi:10.1038/s41598-021-01738-w

[REF:liu2025statin] Liu YH, et al. Gut microbiota profiles by LDL-C target achievement in statin-treated patients: A cross-sectional study. *Nutr Metab Cardiovasc Dis.* 2025;36(3):104465. doi:10.1016/j.numecd.2025.104465

[REF:kamath2024] Kamath S, et al. The atypical antipsychotics lurasidone and olanzapine exert contrasting effects on the gut microbiome and metabolic function of rats. *Br J Pharmacol.* 2024;181(22):4531-4545. doi:10.1111/bph.16507

[REF:qian2023olanzapine] Qian L, et al. Longitudinal Gut Microbiota Dysbiosis Underlies Olanzapine-Induced Weight Gain. *Microbiol Spectr.* 2023;11(4):e0005823. doi:10.1128/spectrum.00058-23

[REF:zeng2024olanzapine] Zeng C, et al. B-GOS alleviates olanzapine-induced lipid disturbances in mice by enriching Akkermansia and upregulation of PGRMC1-Wnt signaling. *Food Chem Toxicol.* 2024;185:114490. doi:10.1016/j.fct.2024.114490

[REF:bertossi2024] Bertossi F. A Possible Role of Akkermansia muciniphila in the Treatment of Olanzapine-Induced Weight Gain. *Cureus.* 2024;16(3):e55733. doi:10.7759/cureus.55733

[REF:dias2024] Dias MF, et al. Effects of antipsychotics on the gastrointestinal microbiota: A systematic review. *Psychiatry Res.* 2024;336:115914. doi:10.1016/j.psychres.2024.115914

[REF:hua2022] Hua X, et al. Longitudinal analysis of the impact of oral contraceptive use on the gut microbiome. *J Med Microbiol.* 2022;71(4). doi:10.1099/jmm.0.001512

[REF:brito2025] Brito J, et al. Hormonal birth control is associated with altered gut microbiota β -diversity in physically active females across the menstrual cycle: a pilot trial. *J Appl Physiol.* 2025;138(3):739-745. doi:10.1152/jappphysiol.00008.2025

[REF:clapp2025] Clapp Organski A, et al. Oral Contraceptives Induce Time- and Intestinal Segment-Dependent Shifts in the Gut Microbiota. *Nutrients.* 2025;17(16):2591. doi:10.3390/nu17162591

[REF:baldi2024] Baldi A, et al. Effects of iron supplements and iron-containing micronutrient powders on the gut microbiome in Bangladeshi infants: a randomized controlled trial. *Nat Commun.* 2024;15(1):8640. doi:10.1038/s41467-024-53013-x

[REF:tian2025iron] Tian C, et al. Intramuscular iron supplementation enhances intestinal barrier function in weaned piglets challenged with enterotoxigenic E. coli. *Front Cell Infect Microbiol.* 2025;15:1553639. doi:10.3389/fcimb.2025.1553639

- [REF:feng2025vitd] Feng X, et al. Vitamin D Deficiency Exacerbates d-Galactose-Induced Intestinal Barrier Dysfunction in a Murine Model of Intestinal Aging. *J Agric Food Chem.* 2025;73(39):24725-24740. doi:10.1021/acs.jafc.5c07085
- [REF:wu2025vitd] Wu CW, et al. Vitamin D3 Improves Hypothalamic-Pituitary-Adrenal Axis Function, Immunological Responses, and Gut Dysbiosis in Sleep Desynchrony. *Brain Behav.* 2025;15(11):e71084. doi:10.1002/brb3.71084
- [REF:cao2025omega3] Cao B, et al. Interaction between dietary omega-3 polyunsaturated fatty acids, obesity and gut microbiota in preclinical models: A systematic review of randomized controlled trials. *Diabetes Obes Metab.* 2025;27(9):4643-4661. doi:10.1111/dom.16535
- [REF:roussel2025fishoil] Roussel C, et al. Human Gut Microbes Produce EPA- and DHA-Derived Oxylipins, but not N-Acyl-Ethanolamines, From Fish Oil. *FASEB J.* 2025;39(12):e70713. doi:10.1096/fj.202500752RR
- [REF:guo2023berberine] Guo HH, et al. Berberine is a potential alternative for metformin with good regulatory effect on lipids in treating metabolic diseases. *Biomed Pharmacother.* 2023;163:114754. doi:10.1016/j.biopha.2023.114754
- [REF:kong2020berberine] Kong WJ, et al. Berberine in the treatment of metabolism-related chronic diseases: A drug cloud (dCloud) effect to target multifactorial disorders. *Pharmacol Ther.* 2020;209:107496. doi:10.1016/j.pharmthera.2020.107496
- [REF:vichvila2020] Vich Vila A, Collij V, Sanna S, et al. Impact of commonly used drugs on the composition and metabolic function of the gut microbiota. *Nat Commun.* 2020;11(1):362. doi:10.1038/s41467-019-14177-z
- [REF:garcia2024] Garcia-Santamarina S, et al. Emergence of community behaviors in the gut microbiota upon drug treatment. *Cell.* 2024;187(22):6346-6357.e20. doi:10.1016/j.cell.2024.08.037
- [REF:zimmermann2019] Zimmermann M, Zimmermann-Kogadeeva M, Wegmann R, Goodman AL. Mapping human microbiome drug metabolism by gut bacteria and their genes. *Nature.* 2019;570(7762):462-467. doi:10.1038/s41586-019-1291-3 (PMID: 31158845)
- [REF:haiser2013] Haiser HJ, Gootenberg DB, Chatman K, Sirasani G, Balskus EP, Turnbaugh PJ. Predicting and manipulating cardiac drug inactivation by the human gut bacterium *Eggerthella lenta*. *Science.* 2013;341(6143):295-298. doi:10.1126/science.1235872 (PMID: 23869020)
- [REF:mainirekdal2019] Maini Rekdal V, Bess EN, Bisanz JE, Turnbaugh PJ, Balskus EP. Discovery and inhibition of an interspecies gut bacterial pathway for Levodopa metabolism. *Science.* 2019;364(6445):eaau6323. doi:10.1126/science.aau6323 (PMID: 31196984)

Chapter 23: Probiotics — Promise, Hype, and Hard Truths

23.1 What “Probiotic” Actually Means — and the Gap Between Definition and Marketplace

Walk into any pharmacy in the developed world and you will find, somewhere between the vitamins and the digestive aids, a shelf of products labelled “probiotic.” Some are capsules. Some are sachets. Some are chilled drinks in small plastic bottles with pictures of smiling intestines on the front. They carry claims that range from the cautiously vague — “supports digestive health” — to the breathtakingly specific — “clinically proven to boost immunity.” They cost anywhere from a few dollars to well over a dollar per capsule. They are purchased, according to market research firms, in quantities that make the global probiotic market worth somewhere between seventy and eighty billion US dollars a year, with projections approaching one hundred and thirty billion by the end of this decade. That is roughly three times the entire global market for influenza vaccines.

And yet if you ask a microbiologist what the word “probiotic” actually means — what it is *supposed* to mean — the answer you get back is considerably more modest than anything the packaging suggests.

In 2014, an expert panel convened by the International Scientific Association for Probiotics and Prebiotics — ISAPP — published a consensus statement in *Nature Reviews Gastroenterology & Hepatology* that attempted to nail down the definition once and for all [REF:hill2014]. The panel, chaired by Colin Hill and including leading figures in microbiology, gastroenterology, and nutrition science, reaffirmed a definition that had first been proposed by a joint FAO/WHO working group in 2001: a probiotic is a preparation of “live microorganisms which, when administered in adequate amounts, confer a health benefit on the host.” Twelve words. Every one of them load-bearing.

Live. Not dead, not heat-killed, not freeze-dried into permanent inactivity. The organisms must be alive when they reach you. A yoghurt that has been pasteurised after fermentation — killing the bacteria that fermented it — is not a probiotic, whatever the label says. A capsule that sat in a hot warehouse for six months and now contains only dead cells is not a probiotic. A 2019 analysis of commercial products found that roughly sixty per cent did not contain the viable organisms, at the counts claimed on the label, by the time they reached the consumer [REF:jackson2019]. The definition says “live.” The marketplace often delivers “were alive once.”

In adequate amounts. The dose matters. A product containing ten million colony-forming units of *Lactobacillus rhamnosus* GG is not the same product as one containing ten billion, and the clinical trials that demonstrated a benefit for that strain used a specific dose — typically ten billion CFU per day in adults. A product that contains the right species but the wrong dose — or a dose that was adequate at the time of manufacture but has dropped by ninety per cent by the time the consumer opens the bottle — does not meet the definition. Most commercial products specify their CFU count “at time of manufacture” rather than “at end of shelf life,” a distinction that the label does not draw attention to and that most consumers do not know to look for.

Confer a health benefit. This is the phrase that separates a probiotic from a fermented food. Yoghurt, kimchi, sauerkraut, and kefir all contain live microorganisms. They may be good for you — we discussed this in Chapter 20. But unless a specific strain in a specific product has been shown,

in controlled human studies, to produce a measurable health outcome, the product does not meet the ISAPP definition of a probiotic. It is a fermented food. The distinction sounds pedantic, but it is the difference between a clinical claim and a culinary tradition.

The Hill panel made one more point that deserves emphasis. The health benefit, they said, must be demonstrated for a *specific strain*, not merely for a species. This is perhaps the single most important sentence in the entire consensus statement, and it is the one the marketplace ignores most comprehensively. Let me explain why it matters.

The strain problem

In microbiology, a *species* is a broad category — like saying “dog.” A *strain* is a specific lineage within that species — like saying “border collie.” Two strains of the same bacterial species can differ in their gene content by five to ten per cent, which is enough to make one of them a harmless commensal and the other a dangerous pathogen. *Escherichia coli* K-12 is a standard laboratory organism that could not hurt you if you swallowed a litre of it. *Escherichia coli* O157:H7 is the strain that causes haemolytic uraemic syndrome and kills children. Same species. Wildly different organisms.

The same principle applies to probiotic bacteria. *Lactobacillus rhamnosus* GG — usually abbreviated LGG — is one of the most studied probiotic strains in the world, with more than thirty years of clinical trial data behind it. It has demonstrated benefits for antibiotic-associated diarrhoea in children, for acute gastroenteritis, and for a handful of other specific conditions. But another strain of *Lactobacillus rhamnosus* — say, one isolated from a different source and given a different designation — may have none of these properties. It may not colonise the same way. It may not produce the same metabolites. It may not interact with the immune system in the same fashion. It is the same species, but it is not the same organism in any functional sense that matters to the patient.

This is not a theoretical concern. The World Gastroenterology Organisation, in its 2024 updated guidelines on probiotics, stated explicitly that recommendations must link specific strains to specific benefits based on human studies, and that systematic reviews that pool results across different strains are valid only if a shared mechanism of action has been demonstrated [REF:wgo2024]. In practice, this means that when a clinical trial demonstrates that *Saccharomyces boulardii* CNCM I-745 reduces antibiotic-associated diarrhoea, that finding applies to *Saccharomyces boulardii* CNCM I-745, sold commercially as Florastor. It does not automatically apply to a different *Saccharomyces boulardii* strain sold in a different capsule by a different company, even if the label says the same species name.

Now go back to that pharmacy shelf. How many of those products specify the strain? Some do — the better ones, the ones backed by companies that fund clinical research on their own products. Many do not. They say “*Lactobacillus acidophilus*” without specifying which *Lactobacillus acidophilus*. They say “probiotic blend” without telling you which blend, in which proportions, at which dose. And they are perfectly legal, because in most countries probiotics are regulated as dietary supplements, not as drugs, and dietary supplement regulations do not require the kind of strain-level identification and dose-level standardisation that drug regulations do.

The gap between the ISAPP definition and the pharmacy shelf is, I think, the single most important thing to understand before reading the rest of this chapter. The definition describes a medical intervention: a specific organism, alive, in a specific dose, with a demonstrated benefit. The marketplace sells a consumer product: a capsule containing some bacteria, probably alive, in some

quantity, with a health claim that may or may not be supported by evidence — and if it is supported, the evidence may have been generated using a different strain, a different dose, and a different population from the one now buying the product. This is not fraud, exactly. It is something more subtle and in some ways more corrosive: a regulatory structure that allows the language of medicine to be used without the rigour of medicine.

Everything that follows in this chapter needs to be read against that background. When I say “probiotics work for condition X,” I mean that one or two specific strains, at specific doses, in specific populations, have been shown to work in controlled trials. When I say “probiotics don’t work for condition Y,” I mean that the trials done so far, using the strains and doses tested, have not demonstrated a benefit — which is not the same as proving that no strain at any dose could ever help. The language of strain specificity is cumbersome, and I will not always write out the full strain designation in every sentence. But the reader should keep it in mind throughout: the word “probiotic” is doing a great deal of work, and much of it is work the evidence does not support.

23.2 The Evidence Hierarchy: Which Strains, for Which Conditions, Actually Work?

If we take the ISAPP definition seriously — specific strains, adequate doses, demonstrated benefits — then the natural next question is: which strains, for which conditions, have actually passed the test?

The honest answer is: fewer than you might expect, but more than zero, and the list is not evenly distributed across the conditions that probiotics are most commonly purchased for.

Where the evidence is strong

The strongest evidence in the probiotic literature — the conditions for which multiple randomised controlled trials, meta-analyses, and in some cases Cochrane systematic reviews converge on a positive finding — clusters around a small number of gastrointestinal conditions. I am going to walk through the main ones, and then I am going to tell you what is *not* on the list.

Antibiotic-associated diarrhoea is the condition for which probiotics are most commonly recommended and most commonly purchased. The logic is intuitive: antibiotics disrupt the gut ecosystem, diarrhoea results, and adding “good bacteria” should help restore the balance. The clinical trial evidence supports this logic — to a point. A 2025 Cochrane review pooling forty-seven randomised trials and more than fifteen thousand participants found that probiotics reduced the risk of *Clostridioides difficile*-associated diarrhoea, with a number needed to treat of approximately sixty-five — meaning sixty-five people need to take the probiotic to prevent one case [REF:esmaeilinezhad2025]. For broader antibiotic-associated diarrhoea (not just *C. difficile*), the risk reduction is more substantial: roughly a thirty-three per cent relative reduction and a nine per cent absolute reduction. The strains with the best evidence for this indication are *Saccharomyces boulardii* CNCM I-745 and *Lactobacillus rhamnosus* GG. For *S. boulardii* specifically, a meta-analysis of twenty-seven trials found a sixty-three per cent reduction in antibiotic-associated diarrhoea risk [REF:mcfarland2024].

This sounds like a clear win, and for the individual patient with a genuine risk of antibiotic-associated diarrhoea — someone on a broad-spectrum antibiotic, particularly clindamycin or a fluoroquinolone, or someone with a history of previous episodes — it probably is. But the numbers

deserve a closer look. A sixty-five-person NNT for *C. difficile* prevention means that for every person who benefits, sixty-four take the probiotic and derive no measurable advantage from it. The absolute risk reduction of nine per cent for broader diarrhoea means that ninety-one out of a hundred people would have had the same outcome without the probiotic. These are not negligible benefits — in a hospital population, preventing even a small number of *C. difficile* cases saves lives and enormous cost. But they are modest benefits, and they are benefits of specific strains at specific doses, not of the generic “probiotic” that most people buy at the pharmacy without knowing which strain is inside.

Acute infectious diarrhoea in children is the second area of strong evidence. LGG and *Saccharomyces boulardii* both shorten the duration of acute gastroenteritis in children by roughly one day, a finding replicated across dozens of trials. The effect is real but modest, and the evidence quality, when assessed by formal grading systems, is rated as low to moderate — not because the trials are badly done, but because the heterogeneity between studies (different strains, different doses, different populations, different definitions of “diarrhoea”) makes it difficult to pool the results cleanly.

Infant colic is a surprising entry on the list, but the evidence is unusually clean. A single strain — *Lactobacillus reuteri* DSM 17938 — has been shown in six randomised trials enrolling more than four hundred infants to reduce crying time in colicky breastfed infants from roughly ninety minutes per day to thirty-five minutes per day at three weeks. The number needed to treat is between four and five, which is extraordinary for any intervention in any field of medicine. The catch — and there is always a catch in probiotic research — is that the benefit has been demonstrated only in exclusively or predominantly breastfed infants. In formula-fed infants, the same strain at the same dose does not appear to work [REF:sung2018]. The reason is not known with certainty, but the most likely explanation is that the breast-milk oligosaccharides that selectively feed *Bifidobacterium* in the infant gut also create an environment in which *L. reuteri* can establish; the formula-fed gut, with its different community structure, does not provide the same niche.

Necrotising enterocolitis in preterm infants is the most dramatic success story in the probiotic literature, and I am going to return to it in detail in section 23.7, because it also illustrates the ecological tensions at the heart of this chapter. For now, the headline: a 2025 meta-analysis of fifty-one randomised trials enrolling nearly twelve thousand very preterm or very low birth weight infants found that specific probiotic combinations — particularly *Bifidobacterium* plus *Lactobacillus* — significantly reduced the incidence of necrotising enterocolitis (the most feared gastrointestinal emergency in neonatal intensive care) and all-cause mortality [REF:chen2025nec]. The ProPrem trial, one of the largest individual studies, showed a fifty-four per cent reduction in NEC using a three-strain combination of *Bifidobacterium longum* subsp. *infantis* BB-02, *Streptococcus thermophilus* TH-4, and *Bifidobacterium animalis* subsp. *lactis* BB-12 [REF:jacobs2013]. This is real, life-saving medicine. It is also medicine for a very specific population — extremely premature infants in neonatal intensive care — and it tells us almost nothing about whether a healthy adult should take a probiotic capsule for “gut health.”

Where the evidence is weak or absent

Now let me tell you what is not on the list, because the conditions that are *not* on the list are, by and large, the conditions for which most probiotic products are actually marketed.

“**Immune boosting**” is probably the single most common claim on probiotic packaging, and it is, in the strict sense, not a clinical outcome at all. It is a marketing phrase. The immune system is not

a muscle that can be “boosted” by swallowing a capsule. It is a complex regulatory network that can be pushed in many directions, some of them helpful and some of them harmful — an “over-boosted” immune system is what causes autoimmune disease. The evidence that specific probiotic strains can modulate specific immune parameters in specific populations is real but limited. One seven-month trial of LGG found a thirty-four per cent reduction in respiratory infections in children attending day care. That is a single trial with a specific strain in a specific population. It is not evidence that the phrase “supports immune health” on a multi-strain capsule bought by a forty-year-old office worker means anything at all.

Weight loss is another area of heavy marketing and thin evidence. A handful of trials using specific strains have shown small reductions in body weight, waist circumference, or body fat percentage — on the order of one to two kilograms over eight to twelve weeks. The results are inconsistent between trials, the effect sizes are smaller than those of any conventional weight-loss intervention, and a 2024 systematic review concluded that “greater homogeneity in procedures is needed in future clinical trials to establish clear recommendations” — which is a polite way of saying the field is a mess [REF:pena2024].

“Gut health” as a general concept — the idea that taking a daily probiotic capsule maintains or improves the health of the digestive system in a person who is not otherwise ill — has, as far as I can find, no evidence from randomised trials in healthy adults. The European Food Safety Authority, which regulates health claims on food products in the European Union, has evaluated hundreds of applications for probiotic health claims over the last fifteen years. It has approved *none*. Zero. Not one probiotic product sold in Europe has met the EFSA standard for a permitted health claim. This does not mean that probiotics are useless — the evidence for specific strains in specific conditions, described above, is genuine. It means that the broad, non-specific claims that drive most consumer purchasing have not met the evidentiary standard that the European regulatory system requires.

Irritable bowel syndrome sits in an uncomfortable middle ground. Some trials show benefit. Some do not. A 2023 meta-analysis in *Gastroenterology* found that probiotics as a class modestly improved overall IBS symptoms, but the evidence quality was rated as low to very low, the effect sizes were small, and the results varied enormously by strain [REF:ford2023]. Nine specific strains or combinations showed some evidence of benefit for specific IBS subtypes, but consistent guidance on which strain to choose for which patient is, as the authors noted, “lacking.” For a condition that affects roughly ten per cent of the adult population in Western countries, the inability to offer a specific, evidence-based probiotic recommendation after two decades of research is a striking gap.

The evidence-to-shelf disconnect

I want to end this section with a number that I think captures the problem better than anything else I have found. A 2019 survey of ninety-three probiotic products available in Canadian pharmacies found that only thirty-three — thirty-five per cent — could be linked to any clinical evidence of efficacy. Of those thirty-three, only eighteen provided a potency guarantee through end of shelf life rather than at time of manufacture. Products with fewer strains were *more* likely to be evidence-based than products with more strains, presumably because the multi-strain blends were formulated for marketing appeal rather than clinical logic [REF:jackson2019].

If you are a consumer standing in front of that pharmacy shelf, the odds are roughly two in three that the product you pick up has no clinical evidence behind it at all. If you happen to pick one of the one-third that does, the odds are roughly even that the viable count inside the capsule matches what the label promises. These are not the odds of a mature, evidence-based therapeutic category. They

are the odds of a marketplace in which the scientific story and the commercial story have diverged to the point where they are barely recognisable as accounts of the same subject.

23.3 The Colonisation Problem — Most Probiotics Don't Establish Residence

There is, beneath the evidence questions I have just described, a more fundamental biological problem that most probiotic marketing quietly ignores. It is this: most probiotic organisms, when you swallow them, do not establish residence in your gut.

They arrive. They pass through. Some of them are detectable in stool samples for a few days or weeks. But they do not, in most people, colonise the mucosal surface of the intestine — the surface that matters, the surface where bacteria interact with the immune system, compete with other organisms, and produce the metabolites that influence host physiology. They are tourists, not immigrants.

This is not a surprise to microbiologists, although it would be a surprise to most consumers. In Chapter 5, I described the principle of colonisation resistance — the ecological observation that a diverse, established microbial community resists invasion by newcomers. The principle was first demonstrated experimentally in 1954, when Bohnhoff, Drake, and Miller showed that streptomycin treatment reduced the infectious dose of *Salmonella* required to colonise the mouse gut by a factor of one hundred thousand [REF:bohnhoff1954]. Remove the established community, and the pathogen walks in. Leave the community intact, and the pathogen cannot gain a foothold. The principle applies equally to “beneficial” organisms introduced from outside. A healthy gut, full of established residents competing for nutrients and attachment sites, is not an easy environment for a newcomer to break into — even a newcomer that would, in other circumstances, be perfectly welcome.

The question of whether probiotics actually colonise the human gut was settled, as definitively as such questions can be settled, by Niv Zmora and colleagues at the Weizmann Institute of Science in Israel, in a paper published in *Cell* in September 2018 [REF:zmora2018].

What Zmora actually found

The Zmora study did something that, remarkably, had not been done before. Previous probiotic studies had measured probiotic presence in *stool* — the end product of the gut — and used this as a proxy for what was happening along the intestinal lining. The problem with this approach is that stool is a mixed sample. It contains bacteria from the lumen (the open centre of the gut tube) and bacteria that have been shed from the mucosal surface, and the two populations are not the same. A probiotic organism can transit through the lumen, be detected in stool, and never touch the mucosa at all.

Zmora's team recruited fifteen healthy volunteers and gave them a commercial multi-strain probiotic preparation containing eleven bacterial strains — a combination of *Lactobacillus*, *Bifidobacterium*, *Lactococcus*, and *Streptococcus* species. This was not a homemade laboratory concoction; it was a real commercial product of the kind millions of people buy. Each volunteer underwent two colonoscopy procedures — one at baseline before starting the probiotic, and one during probiotic consumption — and during each procedure the team collected mucosal biopsies from multiple sites: stomach, duodenum, jejunum, terminal ileum, and descending colon. They sequenced the microbial communities in these biopsies, in the stool, and in the probiotic preparation itself, and they

also performed RNA sequencing on the mucosal tissue to see how the host's own gene expression responded.

The headline finding was striking. The probiotic organisms did appear in stool — you could detect them passing through. But their colonisation of the mucosal surface was highly variable between individuals. Some people were “permissive” — the probiotic strains established on the mucosa in at least some regions of the gut, and the host tissue showed measurable changes in gene expression in response. Other people were “resistant” — the probiotic strains passed through the lumen without establishing on the mucosal surface, and the tissue showed no detectable response.

The distinction between permissive and resistant individuals was not random. The resistant individuals had distinctive baseline immune signatures — their mucosal tissue showed higher expression of genes involved in T-cell activation, innate immune responses, and inflammatory signalling. In plain language: their immune system was already primed to repel newcomers, and the probiotics, despite being “friendly” organisms, were treated as newcomers and turned away.

What stool doesn't tell you

The Zmora paper contained a second finding that was, in some ways, more consequential than the first. The stool samples did *not* predict the mucosal colonisation pattern. You could have two people taking the same probiotic, both showing the same increase in probiotic organisms in their stool, and one of them was permissive (the organisms had colonised the mucosa and were interacting with the host) while the other was resistant (the organisms were merely passing through the lumen without establishing). If all you had was the stool data — which is all that most clinical trials have — you would conclude that both people were “colonised.” The mucosal biopsies told a different story.

This finding puts a large question mark over every probiotic trial that has used stool as its primary readout for colonisation. Most of them, of course, have no choice — colonoscopy is invasive, expensive, and not something you can ask a thousand trial participants to undergo. But it means that the published literature on probiotic colonisation is, to a degree that nobody had previously appreciated, measuring the wrong thing. It is measuring transit, not residence. It is like counting the cars on a motorway and calling them residents of the town they are driving through.

23.4 The Personalisation Problem — Mucosal Colonisation Varies Dramatically Between Individuals

The Zmora finding raises an uncomfortable question for anyone who takes a probiotic or recommends one: is there a way to know, in advance, whether you are a permissive host or a resistant one?

In 2018, the answer was no. The immune signatures that distinguished permissive from resistant individuals in the Zmora cohort were identified retrospectively, by analysing the mucosal biopsies after the fact. No one was going to perform a colonoscopy and a transcriptomic analysis on every patient before recommending a probiotic capsule.

Since then, the field has made tentative progress. A 2025 paper in *PLOS Biology* demonstrated that computational metabolic models could predict, with seventy-five to eighty per cent accuracy, whether a given probiotic species would engraft in a given individual's gut, based on the pre-existing metabolic capabilities of their baseline microbiome [REF:heinken2025]. The approach works by

modelling the nutrient competition between the incoming probiotic and the resident community: if the resident community already occupies the metabolic niches the probiotic needs, engraftment fails; if there are open niches, engraftment succeeds. It is, in essence, a computational version of colonisation resistance — predicting who has room and who does not.

This is promising, but it is a long way from clinical practice. It requires shotgun metagenomic sequencing of the patient's stool, construction of a personalised metabolic model, and a computational prediction — none of which are available in a pharmacy or a GP's office. For now, the practical implication of the Zmora finding is simpler and more uncomfortable: when a doctor recommends a probiotic to a patient, or when a patient buys one for themselves, neither party knows whether the organisms in the capsule will establish in that particular gut. The product might work. It might not. The variation is not random noise in the data. It is a fundamental biological feature of the host-microbe interaction, and it means that the population-level results of clinical trials — “probiotics reduce the risk of X by thirty per cent” — conceal enormous individual variation. For some people in those trials, the probiotic colonised and helped. For others, it passed through and did nothing. The average is positive. The individual outcome is a coin toss.

The analogy I find most useful is organ transplantation. When a surgeon transplants a kidney, nobody expects the organ to function identically in every recipient. Compatibility depends on tissue type, immune status, and the recipient's overall biology. We accept this because we understand that the immune system treats foreign tissue as foreign, and whether it accepts or rejects the graft depends on a complex negotiation between the transplant and the host. A probiotic is not a kidney, but the principle is the same: introducing a foreign organism into an established biological system is not a simple act of addition. It is an ecological negotiation, and the outcome depends on both parties.

23.5 When Probiotics Backfire: The Post-Antibiotic Paradox

We have arrived at the finding this book has been building toward since Chapter 1.

In section 1.4, I told you that “good” and “bad” are misleading labels for microbes — that context and ecology matter more than identity. In section 5.7, I explained colonisation resistance — the principle that a full ecological house is hard to invade. In section 18.6, I described what happens when antibiotics empty the house: the race to recolonise vacant niches, the vulnerability window, the question of who gets there first. I asked you, at the end of that section, to wonder what would happen if you threw new competitors into the mix during that vulnerable window.

Now I can tell you the answer.

In September 2018, in the same issue of *Cell* that carried the Zmora colonisation paper, Jotham Suez, Niv Zmora, and their colleagues at the Elinav laboratory published a companion study that is, in my view, one of the most important papers in the modern microbiome literature [REF:suez2018]. The study asked a question that seems, on its face, to have an obvious answer: after a course of antibiotics, does taking a probiotic help the gut microbiome recover?

The experimental design was elegant. Healthy volunteers received a seven-day course of antibiotics (a combination designed to cause substantial microbiome disruption). They were then randomised into three groups. One group received a standard multi-strain probiotic — the same commercial preparation used in the Zmora study — for four weeks. A second group received autologous faecal

microbiota transplantation — their own pre-antibiotic stool, collected before the experiment began, was returned to them via upper endoscopy and colonoscopy. A third group received no treatment and was left to recover spontaneously.

All three groups were followed with the same invasive mucosal sampling protocol used in the Zmora study: biopsies from multiple gut sites, stool sampling, and host transcriptome analysis, at multiple time points over months.

The paradox

The result was the opposite of what anyone buying a post-antibiotic probiotic would expect.

The autologous FMT group — the people who received their own bacteria back — recovered fastest. Their mucosal microbiome returned to its pre-antibiotic configuration within days. The host tissue transcriptome — the pattern of genes the gut lining was expressing — also returned to baseline rapidly. The restoration was, by every measure the study could apply, quick and nearly complete. This was not surprising. You were giving the gut back the community it had lost. The residents returned to their own house.

The spontaneous recovery group — the people who received no treatment at all — recovered more slowly. Their mucosal microbiome took weeks to months to return toward baseline, and the return was incomplete, consistent with the Palleja and Dethlefsen findings we discussed in Chapter 18. Scars remained. But the trajectory was in the right direction: the community was rebuilding itself, gradually, from the surviving remnants and from re-exposure to environmental sources.

The probiotic group — the people who received the commercial multi-strain preparation — recovered *slowest*. Not just slower than the FMT group, which was expected. Slower than the group that received *no treatment at all*.

Let me say that again, because it is the single most counterintuitive finding in the modern probiotic literature. After antibiotic treatment, taking a standard commercial probiotic *delayed* the recovery of the gut mucosal microbiome compared to doing nothing.

Why it happens

The mechanism, once you see it, makes ecological sense — which is why this book spent five chapters building the reader's ecological intuition before arriving here.

After antibiotics, the gut mucosa is, in the language of the Zmora study, in a “permissive” state. The established residents have been killed or depleted. The niches are open. The colonisation resistance that would normally repel newcomers is gone. This is precisely the state in which a probiotic can colonise most effectively — and in the Suez study, the probiotic strains did colonise, readily and extensively. They expanded across the mucosal surface of the post-antibiotic gut with an efficiency they could never have achieved in an intact ecosystem.

But here is the problem. The probiotic organisms, once established in those empty niches, did exactly what any ecological colonist does: they defended their territory. They consumed the available nutrients. They produced metabolites — including, the Suez team showed, soluble factors secreted by *Lactobacillus* species that *actively inhibited the growth of other bacteria* in vitro. They occupied the ecological space that the native community needed in order to return.

The native organisms — the patient’s own pre-antibiotic commensals, the survivors hiding in mucosal crypts and in biofilm remnants along the gut wall — were trying to recolonise the same niches. But instead of returning to empty territory, they found it occupied. The probiotic strains were not pathogens. They were not causing disease. They were “beneficial” organisms by any conventional measure. But in the specific context of a post-antibiotic gut, they were functioning as *ecological competitors* — squatting in niches that belonged to the native community, slowing the return of the organisms that had co-evolved with that particular host over a lifetime.

Think of it this way. Your gut microbiome after antibiotics is like a house that has been emptied by a flood. The original residents are displaced but not destroyed — some are sheltering nearby, waiting for the water to recede so they can return and rebuild. Now imagine that a well-meaning neighbour, hearing that the house is empty, moves a group of strangers into it. The strangers are perfectly nice people. They keep the house clean. They pay the utilities. But when the original residents try to come home, they find the rooms occupied, the fridge stocked with someone else’s food, and the locks changed. The strangers are not hostile. They are simply *there*, and “there” is where the original residents need to be.

This is what the Suez study showed the probiotic organisms doing in the post-antibiotic gut. They colonised rapidly, they defended their niches competitively, and they delayed the return of the native community — the community that was tailored, by years of co-evolution, to that specific host.

The host transcriptome tells the same story

The Suez team also tracked the host’s own response — the gene expression patterns in the gut lining — and the transcriptomic data reinforced the colonisation findings. In the spontaneous recovery group, the host transcriptome gradually returned toward its pre-antibiotic baseline as the native microbiome rebuilt itself. In the autologous FMT group, the return was fast and nearly complete. In the probiotic group, the host transcriptome remained in a disrupted state for longer than in either of the other two groups. The gut lining was responding to the probiotic organisms as though it were still in a state of ecological disruption — because, at the mucosal level, it was. The native community had not returned, and the signals the gut lining was receiving were not the signals of its usual microbial partners.

23.6 The Competition Framework: “Beneficial” Species as Ecological Disruptors

The Suez finding is not an isolated curiosity. It is the specific instance of a general ecological principle, and once you see the principle, it reframes the entire probiotic conversation.

The principle is this: in a complex microbial ecosystem, the *identity* of an introduced organism matters less than the *context* in which it is introduced. A species that is beneficial in one context — in an intact, diverse community where it fills a functional role — can be disruptive in another context — in a depleted, recovering community where it competes with the organisms that need to return. The same organism, the same dose, the same host, but a different ecological state, and the outcome reverses.

This is not a new idea in ecology. Forest ecologists have known for decades that certain tree species — fast-growing, shade-tolerant, nutrient-hungry — can be either valuable members of a mature forest or destructive invaders of a recovering one, depending on the state of the community they

enter. In restoration ecology, the concept of “competitive displacement” describes exactly what the Suez probiotics did: a species introduced to help the system recover ends up delaying recovery by outcompeting the native organisms that need to re-establish.

Beyond antibiotics: other vulnerable states

The post-antibiotic gut is the most dramatic example, but it is not the only scenario in which this principle applies. Any state in which the normal colonisation resistance of the gut has been reduced creates the same vulnerability.

Neonates are born with an essentially empty gut, and the first weeks of life are a critical window during which the founding community establishes and shapes the immune system (Chapter 3). Introducing a probiotic during this window is introducing an ecological competitor at a time when the native community has not yet formed — and the long-term consequences of displacing native colonisers with commercial strains during the founding period are largely unknown. I will return to this in section 23.7.

Immunocompromised patients — those on chemotherapy, organ transplant recipients, people with advanced HIV, patients in intensive care — have impaired intestinal barrier function, reduced immune surveillance, and often depleted gut communities from prior antibiotic use. In these patients, the colonisation resistance that normally keeps probiotics confined to the gut lumen is weakened, and the organisms can translocate — cross the gut wall and enter the bloodstream. Case reports of *Saccharomyces boulardii* fungaemia in ICU patients with central venous catheters are now well documented, including fatal cases. In 2026, a paper in *Emerging Infectious Diseases* reported two paediatric cases of *Lactobacillus rhamnosus* catheter-related bacteraemia after cardiac surgery, one of which was fatal [REF:cdc2026]. Whole-genome sequencing confirmed that the bloodstream isolate was genetically identical to the probiotic strain the children had been given. The organism was not a pathogen in any conventional sense. It was a “beneficial” bacterium that had been introduced into a host whose barriers could not contain it.

Post-FMT patients represent another vulnerable population, though a more subtle one. After a faecal transplant, the gut is undergoing rapid ecological reorganisation. Introducing a probiotic during this window adds a competitor to an already chaotic landscape, and the Suez data suggest that the effect is likely to be the same as in the post-antibiotic scenario: delayed reconstitution of the target community.

Nutrient competition as the mechanism

The mechanistic question — *how* do probiotic organisms delay recovery — has been partially answered by work that followed the Suez paper. The answer, to a first approximation, is nutrient competition.

In a post-antibiotic gut, the niches are open but the resources are finite. The nutrients available — mucosal glycans, dietary polysaccharides reaching the colon, host-derived mucus — can support a certain total biomass. The probiotic organisms, once established, consume a share of those resources. Every calorie consumed by a probiotic *Lactobacillus* is a calorie not available to a native *Roseburia* or *Faecalibacterium* that is trying to re-establish. The competition is not active warfare — the organisms are not producing toxins or engaging in direct combat (though the Suez team’s finding that *Lactobacillus*-secreted factors inhibit other bacteria suggests some element of chemical interference). It is mostly passive: the probiotic is simply eating the food and sitting in the

seat.

The 2025 metabolic modelling paper I mentioned in section 23.4 confirmed this mechanism computationally. When the models predicted successful probiotic engraftment, they also predicted shifts in short-chain fatty acid production and resource allocation within the community — shifts that, in the post-antibiotic context, came at the expense of the recovering native organisms [REF:heinken2025]. The models could predict, based on baseline community composition and nutrient availability, which individuals would experience delayed recovery after probiotic supplementation. The prediction was not perfect, but it was substantially better than chance, and it was driven entirely by metabolic competition — who gets the food.

23.7 The Infant Probiotic Question — NEC Prevention vs Ecological Disruption

Nowhere in the probiotic literature is the tension between benefit and ecological disruption more acute than in the neonatal intensive care unit.

I described the NEC evidence briefly in section 23.2. Let me now expand on it, because the story is both genuinely uplifting — this is a case where probiotics save lives — and genuinely troubling, for reasons that the NEC literature itself has only recently begun to confront.

The case for

Necrotising enterocolitis is the most common and most devastating gastrointestinal emergency in premature infants. It occurs in roughly five to ten per cent of very low birth weight infants (those born weighing less than 1,500 grams), and its mortality rate, once the disease has progressed to the stage requiring surgery, is between twenty and fifty per cent. The survivors often face lifelong complications, including short bowel syndrome and neurodevelopmental impairment. The pathogenesis is complex and not fully understood, but it involves an interaction between an immature gut barrier, an aberrant immune response, and a dysbiotic microbial community — a community that, in extremely premature infants, has had neither the time nor the ecological input to develop normally.

The evidence that specific probiotic strains can prevent NEC in this population is, by the standards of neonatal medicine, very strong. The 2025 meta-analysis of fifty-one trials found that *Bifidobacterium* and *Lactobacillus* combinations — particularly the three-strain combination used in the ProPrems trial — reduced NEC incidence by roughly fifty per cent and all-cause mortality by a statistically significant margin [REF:chen2025nec]. The ProPrems trial itself showed a reduction from 4.4 per cent to 2.0 per cent — not a large absolute number, but in a neonatal intensive care unit, where every case of NEC is a potential death, the difference is enormous. Many NICUs around the world now use prophylactic probiotics as standard of care for very preterm infants.

The ecological concern

And yet.

The premature infant gut is the most ecologically vulnerable environment in the human body. It is a founding community — not a mature one, not a recovering one, but one that is being assembled for the first time. The organisms that colonise in the first days and weeks of life are not just tenants; they

are the founders, and as we discussed in Chapter 3, the founding community shapes the immune system, the metabolic landscape, and the trajectory of the microbiome for years to come.

When we introduce a probiotic into this environment, we are not adding a supplement to an established community. We are inserting a competitor into a founding event. The probiotic organisms, arriving in a gut that has minimal colonisation resistance, can colonise extensively — as they need to, in order to prevent NEC. But the organisms they displace, or prevent from colonising, are the organisms that would have colonised naturally: species transmitted from the mother’s skin, breast milk, and vaginal flora, species that carry the genetic and metabolic signatures of that particular mother-infant pair.

A 2024 study in *Pediatric Research* examined the microbiome of extremely preterm infants who had received prophylactic probiotics and found that the probiotic strains did indeed colonise effectively — but the overall community structure was significantly different from that of infants who had not received probiotics [REF:patole2024]. The probiotic infants had higher proportions of the administered strains and lower proportions of certain native colonisers. Whether this matters clinically — whether the displacement of native colonisers by probiotic strains during the founding period has consequences that extend beyond the NICU stay — is a question that no study has yet been designed to answer, because the follow-up required (years, perhaps decades) is beyond the scope of any probiotic trial conducted so far.

This is the dilemma at its sharpest. The NEC prevention benefit is immediate, measurable, and life-saving. The ecological cost, if there is one, is deferred, unmeasured, and theoretical. A neonatologist faced with a 1,200-gram infant at high risk of NEC is not going to withhold a probiotic that halves the NEC rate because of a theoretical concern about long-term ecological displacement. Nor should they. The evidence for prevention is strong; the evidence for long-term harm is absent. But the absence of evidence is not evidence of absence, and the ecological logic — that displacing native founders with commercial strains during a critical developmental window might have downstream consequences — is sound even if it has not yet been tested.

The honest conclusion is that NEC prevention with probiotics is one of the genuine triumphs of microbiome medicine, and also that the long-term ecological consequences of routine probiotic administration in the founding period of the human microbiome are unknown. Both of these things are true at the same time, and a mature relationship with the evidence requires holding both of them simultaneously.

23.8 Probiotics and the Resistome: Can Probiotics Expand Antibiotic Resistance?

In Chapter 18, I described how antibiotics do not merely kill bacteria — they select for resistance. The survivors of an antibiotic course carry genes that confer survival, and the post-antibiotic community is enriched in these resistance genes. The collection of all antibiotic resistance genes in the gut microbiome is known, somewhat ominously, as the “resistome,” and its expansion is one of the most consequential side-effects of antibiotic use.

What I did not discuss in Chapter 18 — because it belongs in this chapter — is the finding that probiotic supplementation after antibiotics can make the resistome *worse*.

The Suez resistome signal

The first hint came from the Suez 2018 paper itself. In the supplementary analyses, the team noted that the post-antibiotic gut communities of the probiotic group showed a different resistome profile from those of the spontaneous recovery group. The observation was not the central focus of the paper, and it was noted rather than explored. But it caught the attention of the field.

A subsequent study, published in *Nature Microbiology* in 2021 by members of the same Elinav group, examined the resistome question directly [REF:suez2021resistome]. The team performed deep metagenomic sequencing of stool samples from the original Suez cohorts and tracked the abundance and diversity of antibiotic resistance genes over time. What they found was this: in antibiotic-naïve individuals (people who had not recently taken antibiotics), probiotics had a modest and largely benign effect on the resistome — in some “permissive” colonisers, they actually reduced the abundance of certain resistance genes, presumably by displacing the resistant organisms that carried them.

But in the post-antibiotic context — the scenario in which probiotics are most commonly recommended — the picture was different. Probiotic supplementation was associated with an expansion of the resistome, and the expansion was driven not by the probiotic organisms themselves (which, in most cases, did not carry clinically significant resistance genes) but by *indigenous* bacteria — the patient’s own surviving gut organisms.

The mechanism: ecological disruption as a resistance amplifier

The mechanism is indirect, and it took me a while to understand it, so let me walk through it carefully.

After antibiotics, the surviving gut bacteria include a disproportionate number of organisms carrying resistance genes — that is the basic selection effect of antibiotic treatment. These resistant survivors are initially a minority, held in check by the returning native community as it reconstitutes. In the spontaneous recovery scenario, the native community gradually returns, fills the niches, and drives the resistant minority back to low levels through competitive exclusion. The resistome contracts as the community heals.

In the probiotic scenario, the probiotics occupy the niches instead. The native community cannot return as quickly (this is the Suez finding from section 23.5). The resistant survivors, still present in the community, face less competitive pressure from the native organisms that would normally have outcompeted them. With less competition, the resistant organisms persist at higher levels for longer. Some of them expand into niches that neither the probiotics nor the recovering native organisms have filled. The net effect is that the resistome — the total burden of antibiotic resistance genes in the gut — remains elevated for longer in the probiotic group than in the spontaneous recovery group.

The Elinav team found one particularly striking example. In some probiotic-treated individuals, the post-antibiotic period saw a bloom of indigenous organisms carrying vancomycin resistance genes — genes that confer resistance to one of the last-resort antibiotics used for severe hospital infections. The probiotics did not carry vancomycin resistance themselves. The resistance expansion was a *downstream ecological consequence* of the niche disruption the probiotics caused. By delaying the return of the normal community, the probiotics had inadvertently created a window in which resistant organisms could expand.

What it means

I want to be precise about what this finding means and what it does not mean.

It does not mean that taking a probiotic will make you antibiotic-resistant. The resistance genes in question are carried by commensal organisms in the gut, not by pathogens in the blood or the lungs. The risk to the individual patient, from a single course of post-antibiotic probiotics, is almost certainly negligible.

What it means is something broader and more insidious. If hundreds of millions of people take probiotics after antibiotics — as the current marketing encourages them to do — and if each of those episodes produces a modest, transient expansion of the gut resistome, then the aggregate effect, across populations and across years, is a slow inflation of the resistance gene pool. Resistance genes in commensal organisms are not inert. They can be transferred horizontally, via plasmids and transposons, to pathogenic organisms. The gut is one of the most active sites of horizontal gene transfer in the human body (Chapter 16). Every resistance gene that persists in a gut commensal is a gene that *could*, under the right circumstances, end up in a pathogen.

This is not a reason to panic, and it is not a reason to tell people never to take a probiotic after antibiotics. The NNT of sixty-five for *C. difficile* prevention represents real clinical benefit, and for a patient at genuine risk of *C. difficile*, the benefit almost certainly outweighs the marginal resistome cost. But it is a reason to be thoughtful — to reserve post-antibiotic probiotics for patients who actually have risk factors for antibiotic-associated diarrhoea, rather than recommending them reflexively to everyone who receives a prescription for amoxicillin. The marketing encourages the reflex. The ecology argues for restraint.

A 2025 study of preterm infants added a further nuance. In very low birth weight infants receiving both antibiotics and probiotics, the probiotics actually *reduced* the persistence of certain resistance genes — the opposite of what was seen in adults [REF:reyman2025]. The likely explanation is that in the neonatal gut, where colonisation resistance is minimal and the resistome is dominated by hospital-acquired organisms, the probiotic strains displaced some of the resistant hospital bacteria rather than enabling them. Context, once again, is everything.

23.9 Next-Generation Probiotics: *Akkermansia*, *Faecalibacterium*, Engineered Strains

If the story so far has been cautionary, this section is about where the field is going — and the direction is, despite everything I have just described, genuinely exciting.

The probiotics that most people buy at the pharmacy are, from a scientific standpoint, first-generation products. They are based on organisms — *Lactobacillus*, *Bifidobacterium*, *Saccharomyces* — that were originally isolated from fermented foods or from the human gut decades ago, selected because they were easy to grow, easy to keep alive in a capsule, and generally recognised as safe. The clinical evidence for some of them, as I have described, is real. But their selection was driven by practicality, not by ecological logic, and they are, in many cases, poor matches for the niches they are expected to fill.

The next generation of probiotics is being selected differently — not for ease of manufacture but for ecological fit, and in some cases, engineered for specific therapeutic functions.

***Akkermansia muciniphila*: the mucus-layer sentinel**

Akkermansia muciniphila has appeared several times in this book — in Chapter 20, where its depletion on low-fibre diets was linked to mucus layer thinning, and in Chapter 22, where its enrichment by metformin was identified as a possible mediator of the drug’s metabolic benefits. It is a Gram-negative, obligately anaerobic bacterium that lives in the mucus layer of the colon, where it feeds on host-derived mucin glycans. In a healthy Western adult, it typically represents one to four per cent of the total gut community. It is depleted in obesity, in type 2 diabetes, in inflammatory bowel disease, and in several other metabolic and inflammatory conditions. Its presence is associated, across dozens of observational studies, with leaner body weight, better glucose tolerance, lower systemic inflammation, and a more intact gut barrier.

The step from “associated with health” to “improves health when administered” is the step that most probiotic candidates fail to make. *Akkermansia* appears to be making it.

A 2024 randomised controlled trial published in *Cell Metabolism* gave a pasteurised preparation of *Akkermansia muciniphila* (strain AKK-WST01) to fifty-eight overweight or obese adults with type 2 diabetes for twelve weeks [REF:depommier2024]. The headline finding was positive: the treatment group showed significant reductions in body weight, fat mass, and HbA1c compared to placebo. But the more interesting finding was the subgroup analysis. The benefit was concentrated almost entirely in participants who had *low* baseline levels of *Akkermansia* in their gut. Participants who already had high levels showed no improvement. The organism was helpful precisely when it was filling a gap, and unnecessary when the gap was already filled.

This is the personalisation principle from section 23.4, validated in a randomised trial. It suggests that the future of probiotic medicine is not “everyone takes the same capsule” but “we measure what you’re missing, and give you that.” The concept is straightforward. The implementation — requiring baseline microbiome profiling for every patient — is not.

A separate 2025 trial of a five-strain consortium that included *Akkermansia* showed improvements in postprandial glucose and HbA1c in type 2 diabetes patients over twelve weeks, with no safety signals. And a 130-participant trial of *Akkermansia muciniphila* PROBIO showed significant weight reduction with both viable organisms and with pasteurised (heat-killed) preparations — suggesting that at least some of the benefit comes from bacterial surface molecules rather than from living organisms, which would simplify manufacturing and storage enormously [REF:yoona2025].

A word of caution. A 2024 comprehensive safety review noted that while *Akkermansia* appears safe in healthy and mildly overweight individuals, there are conditions — inflammatory bowel disease, certain autoimmune diseases, post-antibiotic dysbiosis — where its mucin-degrading activity might exacerbate rather than improve mucosal integrity [REF:ashrafian2024]. The organism eats mucus. In a healthy gut with robust mucus production, this is fine — the mucus is continuously replenished, and *Akkermansia*’s activity stimulates further production. In a gut where the mucus layer is already compromised, adding an organism that degrades mucus might not be wise. Context, once again.

***Faecalibacterium prausnitzii*: the butyrate factory**

Faecalibacterium prausnitzii is the single most abundant bacterium in the healthy adult colon, typically representing five to fifteen per cent of the total community. It is the gut’s premier butyrate producer — the organism most responsible for feeding the colonocytes and maintaining the anti-inflammatory tone of the colonic immune system. Its depletion is one of the most consistently

observed features of inflammatory bowel disease, and restoring it has been a goal of microbiome researchers for more than a decade.

The problem has always been practical: *Faecalibacterium prausnitzii* is an extremely strict anaerobe, meaning it dies on contact with oxygen. Keeping it alive in a capsule, shipping it at room temperature, and delivering it to the colon in viable form are all immense technical challenges. Recent work has made progress — microencapsulation techniques and anaerobic formulation strategies are being developed — but a commercial *Faecalibacterium* probiotic suitable for widespread use remains, as of this writing, at the research stage.

What the research shows, however, is tantalising. A 2025 paper in the *European Respiratory Journal* demonstrated that higher intestinal *Faecalibacterium* levels were associated with better clinical outcomes in community-acquired pneumonia and in ICU patients — the first evidence that a gut butyrate producer might influence respiratory immunity [REF:enaud2025]. A 2024 study showed that pre-colonisation with *Faecalibacterium* protected against intestinal barrier damage from sleep deprivation in mice, enhancing goblet cell count, mucin production, and tight junction protein expression [REF:li2024faecal]. The organism is, by every measure, a keystone species of the gut ecosystem — the anchor tenant whose presence keeps the whole building standing.

Whether it can be delivered as a probiotic remains to be seen. But the effort is a signal of where the field is heading: away from generic “probiotic blends” and toward the restoration of specific, ecologically critical organisms that are depleted in specific disease states.

Engineered probiotics: programmable living medicines

The furthest frontier of probiotic science is not in the pharmacy but in the synthetic biology laboratory. Engineered probiotics — bacteria that have been genetically modified to perform specific therapeutic functions — are entering clinical development for a range of conditions that no conventional probiotic could address.

The concept is straightforward, even if the execution is not. Take a well-characterised probiotic chassis — *Escherichia coli* Nissle 1917 is the most common, a non-pathogenic strain that has been used as a probiotic in Germany since 1917 — and equip it with genetic circuits that allow it to sense a disease signal and produce a therapeutic response. The engineered organism colonises the gut, detects the target molecule or environmental condition, and releases its payload locally, at the site where it is needed.

Several of these constructs are now in clinical trials. Engineered bacteria that detect inflammatory markers and release anti-inflammatory cytokines are being tested for inflammatory bowel disease. Others are designed to produce enzymes that break down toxic metabolites in inherited metabolic diseases. A 2026 review in *Theranostics* described engineered probiotics that exploit the hypoxic environment of solid tumours to deliver cytokines, prodrug-converting enzymes, and immune checkpoint inhibitors directly to the tumour site — a form of cancer therapy that uses the bacterium as a guided missile, colonising the tumour and releasing its therapeutic payload from inside [REF:liu2026].

These are not products you will find on a pharmacy shelf any time soon. They are regulated as drugs, not supplements, and they must pass through the full regulatory pathway of clinical trials, safety assessment, and manufacturing controls. But they represent a fundamentally different conception of what a “probiotic” can be: not a capsule of bacteria you swallow for vague digestive benefit, but a

programmable living medicine, designed for a specific disease, with a defined mechanism of action and a measurable therapeutic endpoint.

The irony is striking. The first-generation probiotics were discovered in fermented foods and marketed as supplements, largely outside the framework of clinical medicine. The next-generation probiotics are being engineered in laboratories and developed as drugs, within the strictest regulatory frameworks available. The two generations could hardly be more different in their origins, their evidence base, or their regulatory status. But both call themselves “probiotics,” and the consumer, standing in front of that pharmacy shelf, has no easy way to distinguish between them.

23.10 Regulatory Grey Zones, the Supplement Industry, and What the Consumer Should Know

The regulatory landscape for probiotics is, in most countries, a mess. It is a mess for historical reasons, and understanding those reasons is the shortest path to understanding why the pharmacy shelf looks the way it does.

How we got here

In the United States, probiotics occupy a regulatory no-man’s-land. The Food and Drug Administration does not recognise “probiotic” as a regulatory category. A product containing live bacteria can be classified as a *food* (if it is a yoghurt or a fermented drink), as a *dietary supplement* (if it is a capsule or a powder with no disease claims), or as a *drug* (if it makes a claim to treat, cure, or prevent a specific disease). The classification determines the regulatory burden. A drug must demonstrate safety and efficacy in clinical trials before it can be sold. A dietary supplement must merely be safe and must not make explicit disease claims — but it can make “structure/function” claims (“supports digestive health,” “maintains immune function”) that, to the average consumer, are indistinguishable from disease claims.

The consequence of this framework is that the vast majority of probiotic products are sold as dietary supplements, under rules that do not require clinical trials, do not require strain identification beyond the species level, do not require proof that the organisms are alive at time of purchase, and do not require evidence that the product does what the packaging implies it does. The manufacturer is responsible for ensuring safety. The FDA can act if a product is demonstrably unsafe or if its labelling is fraudulent. But the FDA does not pre-approve supplements, and its enforcement resources are limited. The practical result is that the probiotic supplement market is regulated by the honour system.

In the European Union, the situation is different in detail but similar in outcome. The EU has taken a harder line on health claims — as I mentioned, EFSA has approved zero probiotic health claims to date — but the products themselves are widely available, and in several EU countries the word “probiotic” has been banned from labels precisely because regulators determined that it implies a health benefit that has not been demonstrated. In Italy, by contrast, probiotics are regulated under a specific framework that requires strain identification and minimum viable counts. The regulatory patchwork across the EU is itself a problem for consumers who buy products across borders.

In Australia, the Therapeutic Goods Administration classifies most probiotics as “listed medicines” — a category that requires evidence of safety but sets a lower bar for efficacy than “registered medi-

cines” (which require full clinical trial evidence). The practical consequence is similar to the US: products that make modest health claims can be sold without rigorous evidence.

What the consumer should look for

I am going to close this section with a practical checklist — not because I think a popular-science book should be a shopping guide, but because the reader who has followed this chapter to this point is entitled to know what distinguishes a credible probiotic product from a marketing exercise.

First, *the strain should be identified*. Not just the species — the strain. A product that says “*Lactobacillus rhamnosus*” without specifying the strain designation is a product that cannot be matched to any clinical trial. A product that says “*Lactobacillus rhamnosus* GG” can be checked against the published evidence.

Second, *the CFU count should be guaranteed through end of shelf life*, not at time of manufacture. If the label says “10 billion CFU at time of manufacture,” what you are holding might contain 10 billion, or 5 billion, or 500 million — you have no way of knowing. A product that says “10 billion CFU at time of consumption” or “through expiration date” is making a stronger commitment.

Third, *the claimed benefit should be supported by at least one published randomised controlled trial using the specific strain and dose in the product*. This information is sometimes on the packaging, sometimes on the manufacturer’s website, and sometimes absent entirely. If it is absent, that is your answer.

Fourth, *the product should come from a manufacturer with a track record of clinical research*. Companies that fund and publish clinical trials on their own products — and that register those trials prospectively, before the results are known — are operating in a different universe from companies that formulate a blend of popular species, put it in a capsule, and write “supports gut health” on the label.

Fifth, and finally, *the condition should be one for which probiotic evidence exists*. If you are taking a probiotic because you are on a course of antibiotics and have a history of antibiotic-associated diarrhoea, and the product contains *Saccharomyces boulardii* CNCM I-745 or *Lactobacillus rhamnosus* GG at an adequate dose, you are making a decision supported by evidence. If you are taking a probiotic because you have a vague sense that it is “good for your gut,” the evidence does not support that decision — not because it has been tested and failed, but because it has not been adequately tested at all.

23.11 A Measured Conclusion: Not “Probiotics Are Bad” but “Probiotics Are Drugs We’re Prescribing Without a Diagnosis”

I want to end this chapter by saying clearly what the chapter is not arguing. It is not arguing that probiotics are a scam. It is not arguing that nobody should take them. It is not arguing that the science is all negative. The evidence for specific strains, in specific conditions, at specific doses — antibiotic-associated diarrhoea, NEC prevention in preterm infants, infant colic in breastfed babies — is genuine, replicated, and in some cases life-saving. Anyone who reads this chapter as an argument for throwing away all probiotics has misread it.

What the chapter is arguing is something more nuanced and, I think, more important. It is arguing that probiotics are ecological interventions in a complex microbial ecosystem, and that treating

them as generic health supplements — to be taken by anyone, for anything, without reference to the patient's specific microbiome, the specific strain in the capsule, or the ecological state of the gut at the time of administration — is not just scientifically unsupported but, in some circumstances, counterproductive.

The Suez finding — that probiotics delay post-antibiotic microbiome recovery — is not an argument against probiotics. It is an argument for *context*. In an intact gut, the colonisation resistance that repels probiotics is itself a sign of health, and the probiotic passes through without harm and without benefit. In a post-antibiotic gut, the same probiotic colonises readily but occupies niches needed by the recovering native community. In a premature infant gut, the probiotic may prevent NEC and save a life, while simultaneously displacing native colonisers whose long-term absence we cannot yet measure. The organism does not change. The context changes. The outcome changes.

This is the lesson that ecology teaches about all interventions in complex systems: the same action, in different circumstances, produces different results. A controlled burn restores a prairie. A controlled burn in the wrong season, with the wrong fuel load, destroys it. The burn is the same. The system is different.

The phrase I keep coming back to is one I heard from a microbiologist at a conference several years ago, and I have not been able to improve on it: “Probiotics are drugs we’re prescribing without a diagnosis.” We are giving patients a microbial intervention without knowing what their microbiome looks like, without knowing whether the organisms in the capsule will colonise their particular gut, without knowing whether the ecological state of their gut at that moment makes colonisation helpful or harmful, and without monitoring the outcome at the mucosal level where it matters. We are, in effect, prescribing a treatment for a condition we have not diagnosed, in a patient we have not examined, using a preparation we have not verified, and measuring the result with a method (stool sampling) that does not reliably reflect what is happening at the site of action (the mucosa).

This is not a counsel of despair. It is a description of where the field is right now — in transition between a first generation of probiotics that were developed without ecological thinking and a second generation that is being developed with it. The future, I believe, will look something like this: a patient presents with a condition for which microbiome intervention is appropriate. Their gut microbiome is profiled — cheaply, quickly, perhaps from a stool sample analysed by machine learning. The profile identifies which species are depleted, which niches are open, and which metabolic functions are impaired. A specific organism, or consortium of organisms, is selected to fill the gaps — not a generic blend, but a tailored prescription, chosen for that patient's ecosystem. The intervention is monitored, and adjusted if the organism does not engraft or if the ecosystem responds in an unexpected way.

This is precision microbiome medicine. It does not exist yet, except in research settings. But the tools to build it — shotgun metagenomics, metabolic modelling, personalised prediction algorithms, next-generation organisms selected for ecological fit — are all in active development. The gap between where we are and where we need to be is not a gap of knowledge (we understand the principles) but a gap of implementation (the tools are not yet cheap, fast, or scalable enough for routine clinical use).

In the meantime, the consumer stands in front of the pharmacy shelf, holding a capsule that may or may not contain the organisms the label claims, at a dose that may or may not be adequate, with evidence that may or may not apply to the strain inside, for a condition that may or may not respond to the ecological intervention the capsule represents. The consumer is not stupid. The consumer is underserved — by a regulatory system that does not require the information the consumer needs,

by a marketplace that profits from the information gap, and by a medical profession that has not yet integrated microbiome science into routine clinical decision-making.

The book you are reading is, in part, an attempt to close that gap. The microbiome is not a marketing opportunity. It is an ecosystem — your ecosystem — and it deserves the same respect we would give to any other ecosystem we depend on for survival. You would not fertilise a forest without knowing what soil it grows in. You should not supplement a microbiome without knowing what community it contains.

Key Chapter References

[REF:hill2014] Hill C, Guarner F, Reid G, et al. Expert consensus document. The International Scientific Association for Probiotics and Prebiotics consensus statement on the scope and appropriate use of the term probiotic. *Nat Rev Gastroenterol Hepatol*. 2014;11(8):506-514. doi:10.1038/nrgastro.2014.66

[REF:zmora2018] Zmora N, Zilberman-Schapira G, Suez J, et al. Personalized Gut Mucosal Colonization Resistance to Empiric Probiotics Is Associated with Unique Host and Microbiome Features. *Cell*. 2018;174(6):1388-1405.e21. doi:10.1016/j.cell.2018.08.041

[REF:suez2018] Suez J, Zmora N, Zilberman-Schapira G, et al. Post-Antibiotic Gut Mucosal Microbiome Reconstitution Is Impaired by Probiotics and Improved by Autologous FMT. *Cell*. 2018;174(6):1406-1423.e16. doi:10.1016/j.cell.2018.08.047

[REF:bohnhoff1954] Bohnhoff M, Drake BL, Miller CP. Effect of streptomycin on susceptibility of intestinal tract to experimental Salmonella infection. *Proc Soc Exp Biol Med*. 1954;86(1):132-137. doi:10.3181/00379727-86-21030

[REF:wgo2024] World Gastroenterology Organisation. Global Guidelines: Probiotics and Prebiotics. *J Clin Gastroenterol*. 2024;58(6):498-531. doi:10.1097/MCG.0000000000002002

[REF:esmaeilinezhad2025] Esmaeilinezhad Z, et al. Probiotics for the prevention of Clostridioides difficile-associated diarrhea in adults and children. *Cochrane Database Syst Rev*. 2025;(9):CD006095.pub5. doi:10.1002/14651858.CD006095.pub5

[REF:mcfarland2024] McFarland LV, et al. Can the evidence-based use of probiotics mitigate antibiotic-associated dysbiosis? A systematic review and meta-analysis of *Saccharomyces boulardii*. *Open Forum Infect Dis*. 2024;11(11):ofae615. doi:10.1093/ofid/ofae615

[REF:sung2018] Sung V, et al. Lactobacillus reuteri to treat infant colic: a meta-analysis. *Pediatrics*. 2018;141(1):e20171811. doi:10.1542/peds.2017-1811

[REF:chen2025nec] Chen Y, et al. Effect of probiotics on necrotizing enterocolitis in preterm infants: a network meta-analysis of randomized controlled trials. *BMC Pediatr*. 2025;25:469. doi:10.1186/s12887-025-05469-z

[REF:jacobs2013] Jacobs SE, Tobin JM, Opie GF, et al. Probiotic effects on late-onset sepsis in very preterm infants: a randomized controlled trial (ProPremS). *Pediatrics*. 2013;132(6):1055-1062. doi:10.1542/peds.2013-1339

[REF:ford2023] Ford AC, et al. Efficacy of prebiotics, probiotics, and synbiotics in irritable bowel

syndrome and chronic idiopathic constipation: systematic review and meta-analysis. *Am J Gastroenterol.* 2023;113(10):1547-1548. doi:10.1038/s41395-023-00557-0

[REF:jackson2019] Jackson SA, et al. Improving end-user trust in the quality of commercial probiotic products. *Front Microbiol.* 2019;10:739. doi:10.3389/fmicb.2019.00739

[REF:pena2024] Peña-Jorquera H, et al. Probiotics for weight loss: a systematic review and meta-analysis. *Nutrients.* 2024;16(12):1883. doi:10.3390/nu16121883

[REF:heinken2025] Heinken A, et al. Metabolic modeling reveals determinants of prebiotic and probiotic treatment efficacy across multiple human intervention trials. *PLOS Biol.* 2025;23(5):e3003638. doi:10.1371/journal.pbio.3003638

[REF:suez2021resistome] Suez J, et al. Probiotics impact the antibiotic resistance gene reservoir along the human GI tract in a person-specific and antibiotic-dependent manner. *Nat Microbiol.* 2021;6:1043-1054. doi:10.1038/s41564-021-00920-0

[REF:reyman2025] Reyman M, et al. Impact of early life antibiotic and probiotic treatment on gut microbiome and resistome of very-low-birth-weight preterm infants. *Nat Commun.* 2025;16:4821. doi:10.1038/s41467-025-62584-2

[REF:patole2024] Patole S, et al. Impact of probiotics on gut microbiome of extremely preterm or extremely low birthweight infants. *Pediatr Res.* 2024;96:1178-1186. doi:10.1038/s41390-024-03520-w

[REF:depommier2024] Depommier C, et al. Akkermansia muciniphila supplementation in patients with overweight/obese type 2 diabetes: efficacy depends on its baseline levels in the gut. *Cell Metab.* 2024;37(1):P196-210.E6. doi:10.1016/j.cmet.2024.12.010

[REF:yoons2025] Yoon HS, et al. Akkermansia muciniphila PROBIO ameliorates overweight: a randomized, double-blind, placebo-controlled trial. *Food Sci Hum Wellness.* 2025;14(4):9250659. doi:10.26599/FSHW.2025.9250659

[REF:ashrafian2024] Ashrafian F, et al. Akkermansia muciniphila: a promising probiotic against inflammation and metabolic disorders. *Virulence.* 2024;15(1):2375555. doi:10.1080/21505594.2024.2375555

[REF:enaud2025] Enaud R, et al. Butyrate-producing gut bacterium Faecalibacterium prausnitzii protects against bacterial pneumonia. *Eur Respir J.* 2025;65(7):2208-2024. doi:10.1183/13993003.08-2024

[REF:li2024faecal] Li Y, et al. Faecalibacterium prausnitzii supplementation prevents intestinal barrier injury and gut microflora dysbiosis induced by sleep deprivation. *Nutrients.* 2024;16(8):1100. doi:10.3390/nu16081100

[REF:liu2026] Liu J, et al. Advances in the engineering of living probiotics for cancer immunotherapy. *Theranostics.* 2026;16(3):1164-1183. doi:10.7150/thno.107289

[REF:cdc2026] Zhang L, et al. Case series of bacteremia associated with probiotic use in children after cardiac surgery, China. *Emerg Infect Dis.* 2026;32(1):250298. doi:10.3201/eid3201.250298

[REF:lawley2013] Lawley TD, Walker AW. Intestinal colonization resistance. *Immunology.* 2013;138(1):1-11. doi:10.1111/j.1365-2567.2012.03616.x

Chapter 24: Faecal Microbiota Transplantation — The Nuclear Option

24.1 A Title That Has to Be Earned

The working title of this chapter, in the outline I have been keeping for the last two years, has been *The Nuclear Option*, and I have kept it even as other chapter titles came and went. There is something unavoidably right about the phrase. Everything else in the therapeutic part of this book — probiotics, prebiotics, diet, faecal short-chain fatty acids, phage cocktails, next-generation *Akkermansia* preparations — is, however sophisticated, a nudge. A nudge to an ecosystem that is still more or less the one you grew up with. Faecal microbiota transplantation is not a nudge. It is a wholesale replacement. You take the resident community of one adult gut, you transfer it into another adult gut, and you hope that the one takes root at the expense of the other. In ecology that is not called perturbation; it is called invasion.

Whether the nuclear metaphor flatters the procedure or libels it is something we will have to come back to at the end of the chapter. Nuclear weapons are blunt, collateral, and irreversible; an FMT is precise in its mechanism of action (deliver a whole microbial community to a site where whole microbial communities are supposed to live), and — as we will see — remarkably forgiving in practice. But the metaphor is useful in one specific way: it captures the scale of the intervention, and the seriousness with which it should be contemplated. This is not a probiotic capsule. This is the transfer of an organ.

I want to start with that framing because a great deal of the public conversation about FMT, even in 2026, begins from a misunderstanding that is roughly the opposite of my starting point. Most people, when they first hear about FMT, think either (a) that it is disgusting and so bizarre it cannot possibly work, or (b) that it is a straightforward, almost folk-medical intervention — *you literally just take some of their poo and swallow it, right?* — that any sufficiently motivated person could organise in a weekend. Both of these conceptions are wrong, and they are wrong in a way that matters. FMT, done properly, is a sophisticated, standardised, heavily regulated procedure that bears almost no resemblance to what a desperate patient would do in their kitchen. And FMT, done in the kitchen, is genuinely dangerous, in ways that have killed people. The gap between the two versions of the procedure is the entire subject of this chapter.

The ground I want to cover is this. First, the ancient prehistory and modern rediscovery of the practice — a story that is older and stranger than most readers suspect. Second, the unambiguous success story: *Clostridioides difficile* infection, which is where FMT works with a reliability that is, for a treatment this young, almost without parallel in medicine. Third — and this is the promised *demythification* — what actually happens at a modern FMT clinic, from the moment a donor walks into the office until the moment a recipient walks back out again. There is an entire industry of professional stool banks, IT systems, freezer archives, donor registries, and standard operating procedures sitting between the caricature and the reality, and the reader should know what it looks like. Fourth, the much messier second front: the attempts to extend FMT to inflammatory bowel disease, metabolic syndrome, autism, cancer immunotherapy, and beyond — conditions in which the signal is real but the success rate is not. Fifth, the safety signal that every serious practitioner

of FMT has been watching ever since the first death, and what it taught us about donor screening. Sixth, the phenomenon that makes the rest of the chapter intellectually honest: the Do-It-Yourself FMT community, why it exists, what they actually do, why their self-reported outcomes are so much better than the scientific literature would predict, and the difficult ethical territory that the existence of this community forces the medical profession to traverse. And finally, the question of where this is all going — towards a future in which FMT as we currently know it may turn out to be a transitional technology, superseded by defined consortia, spore capsules, and engineered communities that do the same job without needing a donor at all.

One declaration before I begin. I am an emergency physician, and I do not perform FMT. What I have seen, across many years of practice, are the patients for whom FMT was either life-saving or for whom it was the last thing that would have helped and did not happen — the elderly woman on her seventh course of vancomycin for recurrent *C. difficile*, frail and tired and visibly frightened by the next relapse; the young adult with ulcerative colitis desperate enough to order a stool kit online; the immunocompromised oncology patient whose family asked me, in a corridor, whether the procedure was as disgusting as it sounded. The evidence in this chapter is the same evidence that shapes clinical practice, and I will try to present it with the same rigour I apply elsewhere in this book. But the clinical texture comes from a practitioner who has, repeatedly, been on the downstream end of the decision about whether to refer. The reader should take that for what it is.

24.2 Four Centuries of Yellow Soup, and What We Can Honestly Claim About It

A small claim that one encounters repeatedly in review articles, and occasionally in popular accounts, is that faecal transplantation was “invented in China in the fourth century CE.” The reference, when a reference is given, is to a passage in a Chinese medical text called the *Zhou Hou Bei Ji Fang* — a “Handbook of Emergency Prescriptions” compiled by the alchemist-physician Ge Hong in the early fourth century — describing the use of a preparation called *huanglong tang*, the “yellow dragon soup”, for severe diarrhoea and food poisoning. The preparation was made from the fermented stool of a healthy individual, administered orally to the sick. A later, more famous medical encyclopaedia — Li Shizhen’s *Bencao Gangmu* of 1578 — systematised the practice, listing several stool-derived preparations for gastrointestinal and febrile illness.

I want to be honest about what this evidence can and cannot support. It is genuine: the texts exist, the preparations are described, and the continuity with modern FMT is philosophically real — someone in fourth-century China had the insight that the effluent of a healthy gut could cure the contents of a sick one, and that insight is, at the level of ecology, correct. But the claim that the fourth-century Chinese were “doing FMT” is an overreach. They were doing something that resembles FMT in its logic without sharing any of its epidemiological machinery. There was no donor screening. There was no concept of a specific microbial pathogen to be displaced. There was no way to know whether a batch worked except by watching whether the patient lived. It is a beautiful pre-scientific intuition, and it should be honoured as that; it should not be turned into a press release.

Much closer to the modern idea is a line of veterinary practice that runs through early-modern Europe and into the twentieth century — Italian physicians in the seventeenth century performing *transfaunation* on ruminants (transferring cud from healthy to sick cows), Bedouin herdsmen feeding fresh camel dung to their own families during dysentery outbreaks in the North African deserts,

the latter reportedly relayed to Allied soldiers in the Second World War, who apparently found it worked. These practices are the real throughline from the ancient to the modern period. They tell us that the empirical observation — *healthy gut contents cure sick guts* — has been rediscovered many times across many cultures, by people who had no way of knowing why.

The modern medical literature begins, properly, in 1958, in Denver, Colorado, in the pages of the journal *Surgery*, with a four-page report by a surgeon called Ben Eiseman and his colleagues [REF:eiseman1958]. The context is worth recovering. The antibiotic era had just begun in earnest; penicillin had moved from miracle to mainstay; broad-spectrum antibiotics were being used, for the first time, to treat serious surgical infections. And a new problem had started to appear. Patients who had been treated with broad-spectrum antibiotics — usually for abdominal surgical infection, or for septicaemia — were, in a small but devastating proportion, developing a secondary illness in which the lining of their colon died in patches and replaced itself with a grey-yellow pseudomembrane. The condition was called *pseudomembranous enterocolitis*, and it had a mortality in the published series of the late 1950s of around seventy-five percent.

The cause of pseudomembranous enterocolitis was not known in 1958. (It would take another twenty years to establish that the bug responsible was *Clostridioides difficile*, and that the reason antibiotics predisposed to the disease was that they cleared the commensal species whose presence normally kept *C. difficile* in check.) Eiseman did not know the cause. What he knew was that he had four patients in front of him, desperately ill, failing conventional treatment, and that the logic of the situation pointed in a particular direction. If the disease was some kind of overgrowth of a single organism in a gut that had been “sterilised” by the antibiotics, then perhaps reintroducing a healthy gut community might allow the overgrowth to be displaced. Eiseman obtained stool from healthy donors, prepared it as a retention enema, and administered it to his four patients. Three of the four recovered promptly. One died of an unrelated complication. The paper ran to four pages, with no statistics, no randomisation, no mechanism — just the bare observation that something that looked like it should not work, did.

The paper attracted almost no attention. Eiseman did not pursue the technique. The antibiotic era was in its optimistic phase, and the idea that the solution to an antibiotic-caused illness might lie in the deliberate reintroduction of faecal bacteria did not fit the spirit of the age. Across the next fifty years, isolated clinicians in different countries continued to do the procedure on an ad hoc basis — an Australian gastroenterologist called Thomas Borody, in particular, built a substantial private practice around it from the late 1980s onward, and has sometimes been called the father of modern FMT [REF:paramsothy2017] — but the published literature remained a trickle. A 2011 systematic review, gathering every case report and case series that had ever been published, found a total of 317 patients across all indications and all countries [REF:gough2011]. That was what “the evidence” looked like, as recently as fifteen years ago.

And then, in early 2013, a single paper in the *New England Journal of Medicine* changed everything.

24.3 Amsterdam, 2013: The Trial That Stopped Itself

The paper was by Els van Nood and colleagues, at the Academic Medical Center in Amsterdam, working with the faecal transplant centre at the Leiden University Medical Center [REF:vannood2013]. The disease they were studying was recurrent *Clostridioides difficile*

infection. A word about that disease is needed before the trial can be properly appreciated.

Clostridioides difficile — *C. difficile*, or *C. diff* in clinical shorthand — is an anaerobic spore-forming bacterium that lives, at low density, in the intestines of a small minority of perfectly healthy adults. It is kept in check there by the presence of a diverse surrounding community of commensal organisms that out-compete it for nutrients, occupy the intestinal niche it would otherwise colonise, and metabolise bile acids in a way that prevents the spores from germinating. Remove that community — almost always by giving the patient a course of broad-spectrum antibiotics — and *C. diff* spores that happened to be present in the gut, or are acquired from the hospital environment, germinate into actively growing cells and begin to produce the toxins that are responsible for the disease. The clinical picture is profuse, watery, foul-smelling diarrhoea; sometimes abdominal pain and fever; in severe cases, pseudomembranous colitis, toxic megacolon, perforation, and death. *C. diff* is the most common cause of healthcare-associated infection in the developed world. In the United States alone, on the most recent estimates, it causes around half a million infections and twenty to thirty thousand deaths every year.

The problem that led to the van Nood trial is a specific one: recurrence. When you treat *C. diff* with an antibiotic — oral vancomycin, or more recently fidaxomicin — you can reliably kill the active, growing bacteria. What you cannot kill are the spores. Spores are the dormant form of the organism, and they are exquisitely resistant to antibiotics, to alcohol, to heat, to almost everything except hypochlorite bleach. The spores sit quietly in the gut wall, in crypts, in the surrounding environment. When the antibiotic course is finished and the patient's gut is still in an antibiotic-depleted, low-diversity state, the spores germinate, the cycle begins again, and the diarrhoea returns. About twenty to thirty percent of *C. diff* patients have at least one recurrence. Once you have had one recurrence, your chance of a second is higher again; by the third and fourth recurrence, the chance of yet another is around sixty percent. What ought to have been a single illness becomes a chronic, life-consuming condition. I have looked after patients on their fifth or sixth recurrence who had not been able to leave the house for nine months.

What van Nood and her colleagues did was, for its time, methodologically daring. They took patients with recurrent *C. difficile* infection and randomised them to one of three arms. The first arm was an “initial” short course of vancomycin, followed by bowel lavage and then the infusion of donor faeces — in solution, suspended in saline — through a nasoduodenal tube directly into the small bowel. The second arm was the then-standard treatment: a two-week course of vancomycin, full stop. The third arm was two weeks of vancomycin plus bowel lavage without the faecal infusion. The primary outcome was simply resolution of diarrhoea without relapse at ten weeks.

The trial was stopped early. An interim analysis showed a difference between the arms so large that the data safety monitoring board judged it unethical to continue. In the faecal-infusion arm, thirteen of sixteen patients (81 percent) were cured after a single infusion; all but one of the remaining patients were cured after a second infusion from a different donor. In the vancomycin-alone arm, four of thirteen patients (31 percent) were cured. In the vancomycin-plus-lavage arm, three of thirteen (23 percent). The p-value for the comparison between the infusion arm and either control arm was less than 0.001. The effect size was enormous — a doubling to trebling of cure rates, in a randomised trial, against an active control.

Two things made this paper transformative beyond its immediate clinical finding. First, the authors sequenced the faecal communities of their recipients before and after the procedure. Pre-transplant, the patients' gut communities were low-diversity, dominated by Proteobacteria, with the signature post-antibiotic flattening that had become familiar in microbiome studies. Post-transplant, within

days, the communities came to resemble those of the healthy donors — high diversity, dominated by Bacteroidetes and *Clostridium* clusters IV and XIVa, with the bile-acid-metabolising species that normally keep *C. difficile* in check restored. The transplant had not simply killed the *C. difficile*; it had restored the ecosystem. This was proof of the ecological hypothesis in a way that no previous study had managed.

Second, the trial broke the social taboo that had kept FMT out of mainstream medicine. It was in the *New England Journal of Medicine*. It was randomised. It was done by serious people in a serious institution. Within eighteen months of publication, FMT clinics had opened in most major academic hospitals in North America, Europe, and Australia. By the time the 2017 international consensus guidelines were published, FMT was the recommended therapy for any patient with a second or subsequent recurrence of *C. difficile* infection [REF:cammarota2017].

It is worth dwelling on how rare this kind of result is. In most branches of medicine, new treatments are introduced slowly and cure rates creep up by a few percentage points at a time. FMT did not do that. It went from a handful of case reports to a randomised trial showing a massive effect, and it did so for a condition — recurrent *C. diff* — where the existing treatments were known to be inadequate. The subsequent literature has been remarkably consistent. A long series of randomised trials and large observational cohorts, summarised across multiple meta-analyses, have confirmed cure rates in recurrent *C. difficile* infection of approximately 85 to 95 percent after a single treatment, with most failures responding to a second. For a disease with an established 60 percent third-recurrence rate under best antibiotic therapy, this is extraordinary. In the austere vocabulary of clinical medicine, FMT *works* for *C. difficile*. What has been harder is finding other things it works for, and that is where the rest of this chapter goes. But before we get there, I owe the reader the demystification that the chapter's subtitle promised.

24.4 What Actually Happens: Inside a Modern FMT Clinic

The “yuck factor” that attaches to FMT in the public imagination does most of its work by keeping the procedure unvisualised. Readers who have not thought about it in detail tend to carry a mental picture that runs something like this: a donor produces a sample in the morning, the sample is handed over the counter in a paper bag, the recipient either swallows it or has it put up their rectum by a clinician in a rubber apron, and everyone tries not to think about what has just happened. That picture is almost entirely wrong. It is wrong at every stage. Let me walk through what actually happens at a modern FMT centre, because the reality is both less gross and more interesting than the caricature.

Finding a donor

Start with the donor. The first thing to understand about FMT is that finding a suitable donor is astonishingly difficult. In a typical modern stool bank — OpenBiome in the United States, the Netherlands Donor Feces Bank in Leiden, the Microba bank in Australia, a growing network in China — fewer than three percent of people who volunteer to donate stool end up qualifying [REF:ianiro2020]. Think about that number. Out of every hundred people who walk in expressing interest, ninety-seven are turned away. For comparison, the corresponding figure for blood donation is somewhere around fifty to eighty percent. Stool donation is harder to qualify for than almost any other form of biological donation in medicine.

The reasons become obvious once you consider what the donor must not transmit. The questionnaire and screening protocol at a modern stool bank is typically longer than the protocol for a bone-marrow donor. It asks about: personal and family history of inflammatory bowel disease, irritable bowel syndrome, autoimmune disease, allergy, atopy, autism spectrum disorder, depression, anxiety, obesity, metabolic syndrome, cancer, chronic fatigue syndrome, and any psychiatric disease; history of antibiotic use in the preceding six months; history of travel to regions with endemic parasitic or bacterial infections; history of any sexual behaviour carrying an elevated risk of bloodborne or gastrointestinal infection; and a full medication list, because many medications — the proton pump inhibitors of the previous chapter, metformin, antidepressants, antibiotics, NSAIDs — substantially reshape the gut community in ways that may not be desirable to transplant. The reason for the exhaustiveness is the growing recognition, in several independent studies across the last decade, that the microbiome travels with the disease. Transplant an “obesogenic” community into a lean recipient, and there is at least a theoretical possibility — and, in one famous case report, a documented actual possibility — that the recipient might gain weight they would not otherwise have gained [REF:alang2015]. Transplant from a depressed donor, and there is ongoing concern (not yet definitively demonstrated) that behavioural traits might travel with the microbial community. The field has not yet decided how seriously to take these risks; while that decision is being made, the donor screens are being kept conservative.

Laboratory screening, once the questionnaire is cleared, is extensive. The candidate donor’s blood is tested for HIV, hepatitis A, B, C, and E, syphilis, HTLV-I and II, and — since 2020 — for SARS-CoV-2 where relevant. Their stool is tested for enteric pathogens (*Salmonella*, *Shigella*, *Campylobacter*, *Yersinia*, *Listeria*, shiga-toxin-producing *E. coli*, *Vibrio*), for parasites (*Giardia*, *Cryptosporidium*, *Entamoeba histolytica*, helminth ova), for *C. difficile* itself, for noroviruses and rotaviruses, and — critically, and this is the lesson of a death we will come to in a moment — for antibiotic-resistant bacteria, particularly extended-spectrum beta-lactamase (ESBL) and carbapenemase-producing *Enterobacteriaceae*, vancomycin-resistant enterococci (VRE), and methicillin-resistant *Staphylococcus aureus* (MRSA) if there is mucosal carriage. The list has lengthened over the past decade as safety incidents have taught the field what it had previously been prepared to overlook.

A prospective donor who clears the questionnaire and the first round of stool and blood screening is then enrolled, which in most modern banks involves a commitment to donate regularly over several months. The bank will want to sequence the donor’s stool microbiome to characterise it, will often retest stool samples at the time of donation, and will quarantine the donated material for up to four weeks before releasing it for clinical use, during which time the donor is re-tested for the same pathogens to catch anything that was in the incubation period at the time of donation. It is the same quarantine-and-re-test logic that is used for whole-blood donation in the setting of potential early HIV or hepatitis transmission.

Processing and storage

The material itself is not what a lay reader might imagine. A donated stool sample arrives at the bank, usually within a few hours of being produced. It is weighed, logged, and moved into an anaerobic processing cabinet — a sealed glove-box in which the atmosphere is oxygen-depleted, because most of the organisms that matter in the gut are obligate anaerobes that begin to die on exposure to air within minutes. Inside the cabinet the stool is homogenised in buffered saline (typically in a ratio of around one to five by weight), filtered through progressively finer sieves to remove undigested food fibre and particulate debris, and cryoprotected with glycerol — usually to a final concentration of about ten percent. The resulting suspension is aliquoted into dose-sized units of

anything from thirty to two hundred millilitres, depending on the intended route of administration, labelled with donor and batch identifiers, and frozen to minus eighty degrees Celsius. There it sits, indistinguishable in appearance and in handling to blood-bank cryoprecipitate, until a clinician somewhere orders it.

A modern stool bank is an industrial operation. OpenBiome, the non-profit stool bank that supplies most of the FMT carried out in the United States, has processed and distributed over sixty thousand treatments since its founding in 2012. It operates with the standard quality-management systems that any pharmaceutical-grade biological manufacturer uses: good manufacturing practice, deviation logs, batch records, chain-of-custody documentation. The material it ships looks, to the clinician who opens the box, like any other frozen biological: a small labelled bottle in a cryo-transport container, accompanied by a paper trail that includes the donor's screening results (de-identified), the processing batch, and the release certificate. Nothing about it is improvised. Nothing about it is handed across a counter in a bag.

Delivery to the recipient

When the material reaches the patient, it can be delivered by any of four established routes, each with specific advantages. The first, and historically the commonest, is *colonoscopic infusion* — the stool suspension is delivered directly into the caecum and ascending colon through the working channel of a colonoscope, after the colon has been prepared (the same bowel preparation as for any other colonoscopy). This has the advantage of depositing the new community at the proximal end of the colon, where it can seed downstream, and of allowing the operator to inspect the bowel mucosa at the same time. It has the disadvantage of requiring a colonoscopy, with the usual preparation, sedation, and small procedural risks. The second route is *retention enema* — the suspension is delivered into the distal colon via a rectal tube, and the patient is asked to retain it for as long as possible. This is simpler and cheaper, and can be done in a day-surgery setting without sedation, but it delivers the new community only to the distal colon; seeding of the ascending colon is less reliable. The third route is *nasojejunal or nasoduodenal infusion* — a tube is passed through the nose, down the oesophagus and through the stomach into the small bowel, and the suspension is delivered from above. This was the route used in the van Nood trial. It has the advantage of seeding the small bowel as well as the colon; it has the disadvantage of being uncomfortable for the patient (anyone who has had a nasogastric tube can confirm this) and of the very small theoretical risk of the new microbial community being aspirated into the lungs.

The fourth route — and this is the one that has reshaped the field in the last decade — is *oral capsules*. The processed, anaerobically-prepared, cryopreserved stool suspension is dispensed into double-walled, acid-resistant capsules designed to survive the stomach acid and release their contents only after they reach the small bowel. The capsules are frozen, shipped on dry ice, and swallowed by the patient like ordinary pills. A typical capsule dose for recurrent *C. difficile* is twenty to forty capsules, split over one or two days. In 2017, Dina Kao and her colleagues at the University of Alberta published a randomised controlled trial directly comparing oral capsule FMT with colonoscopy-delivered FMT for recurrent *C. difficile* [REF:kao2017]. The cure rate in both arms was 96.2 percent. They were identical. And the capsule patients, unsurprisingly, rated their experience as considerably less unpleasant than the colonoscopy patients.

This last result — that capsules work as well as colonoscopy — is what most thoroughly demolishes the caricature. For the commonest indication, the patient's actual experience of FMT in 2026 is essentially: swallow twenty to forty capsules of what looks like ordinary medication, with a glass of

water, over one to two days. There is no yellow liquid. There is no tube. There is no colonoscopy. The capsules have no taste and no smell; they are double-encapsulated precisely to ensure that nothing breaches the outer shell until after the stomach. The visceral horror-film image of FMT — the milkshake, the enema, the bucket, the kitchen blender — is for most modern patients simply not what happens. The procedure has been, in the most literal sense, pharmaceuticalised.

Recovery

Follow-up is remarkable for how little there is to follow up. Most patients report mild abdominal cramping and loose stools for one to three days after the procedure, which is approximately what one would expect from introducing several hundred grams of foreign gut contents to the bowel. The new community establishes within days. In *C. difficile* patients, diarrhoea resolves within a week in the overwhelming majority. In most cases the patient simply feels *normal* again, often for the first time in many months. The long-term follow-up data — which now extend out to more than a decade for some of the earliest patients — show persistence of the transplanted community in a substantial minority, partial engraftment (a hybrid of donor and recipient communities) in many, and, interestingly, some patients who revert to something that looks like their original community but without *C. difficile* disease. What precisely constitutes “success” at the community level is still being worked out. What constitutes success at the patient level — no more diarrhoea, no more fear of the next recurrence, a life returned — is clear enough.

24.5 Beyond *C. difficile*: The Frontier That Does Not Quite Work Yet

If FMT were only a treatment for recurrent *C. difficile*, it would still be one of the most striking additions to the modern pharmacopoeia. But the logic of the procedure — transfer a diverse community from a healthy donor to restore a dysbiotic gut — is so attractive that researchers have, unsurprisingly, extended it to a long list of diseases in which the gut microbiome is implicated. Some of these extensions have shown real, reproducible signals. Others have not. A few have quietly failed. I want to walk through the main ones, because the reader who has heard that “FMT can treat X” should know what the evidence actually looks like.

Inflammatory bowel disease

The most active FMT research front outside of *C. difficile* has been inflammatory bowel disease, particularly ulcerative colitis. The rationale is strong. Ulcerative colitis is characterised by a chronic, relapsing inflammation of the colonic mucosa, and the colonic mucosa is where the microbiome lives most densely. Affected patients have a recognisably altered gut community, with reduced diversity, loss of butyrate-producing organisms (*Faecalibacterium prausnitzii* and *Roseburia*), and expansion of Proteobacteria and other pro-inflammatory taxa. If you are going to test the hypothesis that microbiome manipulation can alter an inflammatory disease of the gut, ulcerative colitis is where you would start.

The landmark randomised trial in this space was published in *The Lancet* in 2017 by Sudarshan Paramsothy and colleagues in Sydney [REF:paramsothy2017]. They randomised eighty-five patients with active ulcerative colitis to receive either intensive-dosing FMT — an initial colonoscopic infusion followed by enemas five days a week for eight weeks, each enema prepared from a pool of faeces from three to seven unrelated donors — or placebo. The primary outcome was steroid-free

clinical and endoscopic remission at week eight. It was met in 27 percent of the FMT group and 8 percent of the placebo group (relative risk 3.6, 95 percent confidence interval 1.1 to 11.9, $p = 0.021$). The result was real. It was also modest. Almost three-quarters of treated patients did not remit, and the effect size, while statistically significant, was nothing like what FMT delivers in *C. difficile*.

What Paramsothy's trial and several subsequent studies have taught is that the extension from *C. difficile* to ulcerative colitis is qualitatively different. In *C. difficile*, the disease is the dysbiosis, and restoring the community resolves it. In ulcerative colitis, the dysbiosis is a component of a multi-factorial disease — genetic, immunological, environmental — and restoring the community modifies but does not cure the underlying process. FMT in ulcerative colitis is a real therapy, but it is a partial one. The most recent meta-analyses put the clinical remission rate at around 25 to 30 percent, against 5 to 10 percent with placebo — a useful but not transformative effect. Current practice in most jurisdictions is to reserve FMT for patients who have failed multiple conventional therapies, and even then only in specialist centres. The procedural protocols that seem to work are the intensive multi-donor ones; single-donor, single-dose protocols have been consistently less effective.

Crohn's disease, the other major form of inflammatory bowel disease, has been much more difficult. The transmural inflammation of Crohn's, the propensity to fistula formation, and the small-bowel involvement make the microbial ecology considerably more complex. The handful of trials that have been done have shown mixed results at best, and most FMT centres do not currently treat Crohn's disease outside of research protocols.

Metabolic syndrome, obesity, and type 2 diabetes

A second major front has been the metabolic diseases. The opening shot was fired by Anne Vrieze and colleagues (including Max Nieuwdorp, who would go on to build much of the field) in 2012, before the van Nood paper that made their laboratory famous [REF:vrieze2012]. Vrieze took men with metabolic syndrome — overweight, insulin-resistant, elevated waist circumference, abnormal lipid profile — and randomised them to receive a small-bowel infusion of faecal microbiota from either a lean, insulin-sensitive male donor or their own pre-collected stool (autologous control). Six weeks after infusion, the lean-donor recipients had improved insulin sensitivity — measured by the gold-standard hyperinsulinemic clamp — by approximately seventy percent. The autologous controls did not. The difference was statistically significant.

This was an astonishing finding. It suggested that metabolic disease could be modified, in adult humans, by the transplantation of a community from a lean donor into an obese recipient. The implications, if confirmed, were enormous. What has happened subsequently is that the original signal has been partly confirmed and partly eroded. Longer follow-up studies have shown that the insulin-sensitivity benefit wanes over months; it is not a durable cure. Attempts to translate the finding into actual weight loss have mostly failed. The most honest current reading is that microbiome manipulation can modulate insulin sensitivity in metabolic syndrome for a period of weeks to months, but it does not, on its own, reverse the underlying metabolic disease. The donor community shifts the recipient's physiology; the recipient's physiology, over time, shifts the donor community back. The ecosystem has equilibria, and they reassert themselves.

The Alang and Kelly case report of 2015 [REF:alang2015] is worth lingering on, because it is the mirror image of the Vrieze result and has shaped FMT donor-screening policy ever since. The case was of a thirty-two-year-old woman with recurrent *C. difficile* infection who underwent FMT from her overweight adolescent daughter. The *C. difficile* was cured. Over the following sixteen months,

the patient gained approximately fifteen kilograms of weight, taking her from normal-weight to frankly obese, despite no reported change in diet or lifestyle. She subsequently met the full criteria for metabolic syndrome, which she had not previously had. The case was a single report, without definitive proof that the weight gain was caused by the transplant rather than by some other factor — but the timing, the magnitude, and the clinical trajectory were striking enough that the case became a landmark in the field. From that point onwards, every major stool bank required donors to be not only healthy but of healthy body-mass-index, and to meet explicit metabolic criteria that excluded the kind of community Alang's daughter had donated.

Autism spectrum disorder

The application of FMT to autism spectrum disorder is among the most contested of the extensions, and deserves careful handling. The rationale derives from the observation, now widely replicated, that autistic children with gastrointestinal symptoms have characteristically altered gut microbiomes, and that the severity of GI symptoms in these children correlates with the severity of behavioural symptoms. The question has been whether correcting the microbiome changes the behaviour.

The most cited work in the area comes from the group of Rosa Krajmalnik-Brown and James Adams at Arizona State University, with Thomas Borody and Alexander Khoruts as clinical collaborators. In 2017 they published an open-label trial of what they called Microbiota Transfer Therapy — a multi-week antibiotic course followed by bowel cleansing and an extended FMT protocol — in eighteen autistic children with gastrointestinal symptoms [REF:kang2017]. At the end of treatment, GI symptoms had improved by approximately eighty percent, and autism-related behavioural symptoms had improved by a clinically meaningful margin. A two-year follow-up, published in 2019, showed that most of the improvements had been maintained, and that the behavioural benefits had in some cases continued to accrue over time [REF:kang2019].

Four caveats, all of which the authors themselves are explicit about, attach to this literature. The study was open-label: both the families and the assessors knew that treatment had been given. The sample size was small. There was no placebo control. And the intervention was complex — antibiotics, bowel cleanse, and FMT together — so the contribution of the FMT alone could not be isolated. Since 2019, a small number of larger, better-controlled trials have been reported, with more mixed results. A properly randomised, placebo-controlled trial in autistic children large enough to settle the question has not yet been published. Several are in progress. My own reading of the current evidence is that there is a real but variable signal in a subset of autistic children with GI symptoms, that the effect is probably mediated by improvement in the GI symptoms (which themselves can be behaviourally disruptive) rather than by a direct effect on core autism features, and that parents considering FMT for their autistic child should do so only within a research protocol at a specialist centre. There is, however, an active DIY-FMT community in the autism space, about which more shortly.

Cancer immunotherapy — the newest and most surprising indication

The most recent, and arguably the most interesting, extension is not a treatment for a microbiome disease at all. It is a use of FMT to potentiate cancer immunotherapy. Immune checkpoint inhibitors — the class of drugs that includes pembrolizumab, nivolumab, ipilimumab, and their congeners — work by taking the brakes off a patient's own T-cell response to tumour. They have revolutionised the treatment of melanoma, non-small-cell lung cancer, and several other malignancies. But they only work in a subset of patients, around a third to a half depending on the cancer type, and the

reasons for non-response have been one of the great unsolved puzzles of oncology.

A series of observational studies in the late 2010s established that the composition of the gut microbiome correlated with checkpoint-inhibitor response, with particular taxa (*Akkermansia muciniphila*, *Faecalibacterium*, *Bifidobacterium*) enriched in responders. The question was whether the correlation was causal — whether you could make a non-responder respond by giving them the community of a responder. Two 2021 papers, published back-to-back in *Science*, attempted to answer that question. In the more widely cited of the two, Diwakar Davar and colleagues at the University of Pittsburgh took fifteen patients with PD-1-refractory melanoma — that is, patients whose tumours had progressed on checkpoint inhibitor therapy — and treated them with FMT from patients whose own melanomas had responded well to PD-1 blockade, followed by a re-challenge with the checkpoint inhibitor [REF:davar2021]. Six of the fifteen — 40 percent — responded. In patients whose prior best response had been progressive disease on the same drug, a 40 percent response rate to FMT-plus-drug is, clinically, a startling finding. The accompanying mechanistic work showed that the responders had acquired immune changes consistent with those seen in natural responders — increased CD8 T-cell activation, decreased interleukin-8-producing myeloid cells — in a way that tracked with the engraftment of donor microbial taxa.

The immunotherapy story is still unfolding. Larger trials are underway, using both FMT and defined consortia derived from super-donor material, and the early signals continue to be positive. If the effect holds up at scale — which is not yet certain — FMT may end up with a substantial and entirely unexpected second indication: not as a gut therapy, but as an immunomodulator in oncology.

Other indications, briefly

I will list, without detail, the other extensions that are in active research in 2026: hepatic encephalopathy (multiple small trials, real effect, unclear mechanism — possibly reduction of ammonia-producing bacteria); irritable bowel syndrome (mixed trials, probably a modest effect in a subset of patients, no clear predictor of who benefits); primary sclerosing cholangitis (early studies, intriguing findings around bile-acid metabolism); chronic fatigue / ME (enthusiastic DIY community, sparse clinical data); Parkinson's disease (small trials, early signals); type 1 diabetes [REF:degroot2020] (one randomised trial in newly-diagnosed patients, with a signal for preservation of beta-cell function). None of these has anything like the evidence base of *C. difficile*. All of them are under active investigation. Some will eventually join *C. difficile* as established indications; many will not.

24.6 The Concept of the Super-Donor

Before leaving the question of efficacy, I owe the reader a concept that has quietly reorganised thinking in the field over the last several years: the super-donor.

The observation that started the conversation was uncomfortable for researchers used to randomised trials. When outcomes from early FMT trials — particularly for non-*C. difficile* indications — were broken down by donor, it became clear that some donors produced vastly better outcomes than others. A donor at one centre might have a 60 percent cure rate in ulcerative colitis; another donor at the same centre, working under the same protocol, might have a 10 percent cure rate. This

“donor effect” turned out to be reproducible across studies. Certain donors simply produced more therapeutically effective material than others, for reasons that were not initially clear.

Over time, the hypothesis has crystallised. The super-donors are, in general, donors whose stool communities combine three features: exceptionally high alpha-diversity (a large number of different taxa present); exceptionally high abundance of beneficial keystone species (*Faecalibacterium prausnitzii*, *Akkermansia muciniphila*, specific butyrate-producers); and — this is the subtler point — a resistance profile that has not been substantially damaged by past antibiotic exposure. The proposed mechanism is not that super-donor stool contains some magical unknown organism; it is that super-donor stool contains a more complete and functionally redundant community, and that such communities are more likely to successfully engraft and to produce the metabolic outputs (particularly short-chain fatty acids from fibre fermentation) that mediate the therapeutic benefit.

The super-donor phenomenon has had two practical consequences. The first is that modern stool banks now stratify their donor pools by microbiome characteristics, not just by disease screening, and match donors to recipients based on which community looks most likely to help. The second is that it has accelerated the move towards multi-donor pooled preparations — the Paramsothy ulcerative colitis trial, for example, pooled three to seven donors per dose, on the theory that increased microbial diversity across donors would buffer against the specific deficiencies of any one donor community. Pooled multi-donor preparations have become standard for most non-*C. difficile* indications. They also, incidentally, complicate the super-donor story: if the pool is doing the work, identifying any one donor’s contribution becomes correspondingly harder.

The super-donor concept is also a pointer to the future. If we can identify, by sequencing, what makes a super-donor super, we can in principle reconstruct the relevant community from defined components — this is what the “defined consortium” companies are trying to do — and move from the unpredictable biological product (one person’s stool) to a manufactured product (a known mixture of known organisms at known doses). Two such products now have regulatory approval for recurrent *C. difficile*, and we will turn to them in a moment. The super-donor is the bridge between FMT-as-it-is-now and FMT-as-it-will-be.

24.7 The Safety Signal We Have to Take Seriously

I have been broadly positive about FMT so far, because the broad reading of the evidence is broadly positive. But the chapter would be dishonest if it did not give equal weight to the single clearest safety signal in the history of the procedure.

On 13 June 2019, the U.S. Food and Drug Administration issued an urgent safety alert describing two cases in which immunocompromised patients had received FMT from a shared donor and had subsequently developed bacteraemia — bloodstream infection — with an extended-spectrum beta-lactamase-producing strain of *Escherichia coli*. One of the patients died. The corresponding full case report, by Zachariah DeFilipp and colleagues at Massachusetts General Hospital, was published in *The New England Journal of Medicine* later that year [REF:defilipp2019]. Genomic sequencing of the *E. coli* recovered from the two patients matched, at high resolution, the strain isolated from the donor’s stool — a strain that had not been screened for at the time of donation, because routine donor screening protocols at the time did not include ESBL-producing Enterobacteriaceae.

This incident matters in several ways. First, at the human level — a patient died who should not

have died, from an intervention intended to help them. That fact must not be abstracted into statistics. Second, the case forced a recognition that had been latent in the field: antibiotic-resistant organisms are widely carried, asymptotically, in the guts of otherwise healthy people in the developed world. Community carriage rates of ESBL-producing *E. coli* are now estimated at five to ten percent in many populations, with higher rates in people who have travelled to regions where such organisms are endemic. A donor whose stool is entirely healthy for their own immune system is not necessarily safe for an immunocompromised recipient. The lesson led directly to a hardening of donor screening: ESBL and carbapenemase-producing *Enterobacteriaceae* are now routinely tested for at all major stool banks, as are VRE and — in some protocols — MRSA.

Third, and more subtly, the incident changed the way the field talks about FMT risk. Before 2019, the prevailing narrative — which I probably shared — was that FMT was a remarkably safe procedure, with adverse events essentially limited to the mechanical risks of the route of administration (colonoscopy or nasogastric intubation) and a small risk of minor GI symptoms. That narrative was not exactly wrong, but it was incomplete. The DeFilipp case demonstrated that FMT has a distinct and specific risk profile — infection transmission from donor to recipient — that cannot be eliminated, only minimised through careful screening. Subsequent analyses have estimated the absolute risk of clinically significant infection transmission with modern screening at somewhere between 0.1 and 1 per 1000 procedures, depending on the population and the organism, but it is not zero and it probably never will be.

Beyond infection transmission, three other classes of long-term safety signal have been watched. The first is the possibility of transmission of *non-infectious* traits — obesity, allergy, autoimmunity, perhaps behavioural features — from donor to recipient. The Alang-Kelly weight-gain case is the clearest single example; other signals exist but remain contested. The second is the unknown long-term consequence of colonising an adult gut with a radically different microbial community. The long-term follow-up studies now extending past ten years have been reassuring; no new syndromes have emerged. But ten years is not a lifetime, and FMT recipients are heterogeneous enough that subtle long-term effects could still be escaping detection. The third is the pandemic-era risk that faecal material can transmit respiratory viruses — SARS-CoV-2 was eventually shown to be shed in stool in a proportion of infected individuals, and stool banks had to add SARS-CoV-2 PCR to their donor screens during 2020 [REF:ianiro2020covid]. Future emerging pathogens will predictably require similar additions to the screen.

For the overwhelming majority of patients receiving FMT for recurrent *C. difficile* — which remains the indication in which the procedure's benefit is largest and best established — the risks are small and the benefits are, in the classical sense, life-saving. For patients with non-immunocompromised, non-urgent indications, the balance is less clear, and the conservative answer in 2026 is that such patients should be treated only within clinical trials at established centres. The DeFilipp case is a reminder that “broadly positive” evidence is not the same as “risk-free”, and that the yellow bottle in the cryo-transport container, despite its pharmaceutical packaging, is still a biological preparation derived from another human being's gut.

24.8 The Underground: DIY-FMT and What It Tells Us

I now come to the section of the chapter that was the user's stated interest when the chapter was commissioned, and that I regard as, in certain respects, the most revealing part of the FMT story as a whole. There is a substantial, persistent, global community of patients who perform FMT on

themselves, without medical supervision, typically in their own kitchens, using stool from a donor they recruit privately — a spouse, a sibling, a friend’s healthy child, sometimes a donor they have found through one of the online patient-support networks. They are doing this in 2026 at rates that, on the best available surveys, continue to rise year on year. This community, usually called the *Do-It-Yourself FMT* or *DIY-FMT* community, is not a fringe or a conspiracy. It is a response — rational in its own terms, dangerous in its methods — to a genuine gap in the medical system, and the reader should understand why it exists, what its members actually do, and what the ethical position of the medical profession ought to be in relation to it.

Why it exists

The community exists because, for many of the conditions that patients believe FMT might help them with, *the medical system will not offer it to them*. In most countries in 2026, FMT is regulated as either a biological therapy or an investigational new drug, which means it is legally available only through a licensed clinician for a specifically approved indication — and the only indication for which it is approved in most jurisdictions is recurrent *Clostridioides difficile* infection. A person with ulcerative colitis who has failed multiple biological therapies, has read the Paramsothy paper, and wants to try FMT, in most jurisdictions, cannot. A person with autistic symptoms and gastrointestinal disease who has read the Kang papers, cannot. A person with chronic fatigue syndrome, with irritable bowel syndrome, with treatment-resistant anxiety, with any of the conditions for which the scientific literature shows a modest signal and clinical practice does not yet reach, cannot. The medical system is, in these cases, holding the line against premature adoption of an unproven therapy — and that is, on a population-health view, often the right position. But from the perspective of an individual patient with a serious chronic illness that has failed conventional treatment, the medical system’s caution looks like a door that has been closed in their face. Behind the door is a therapy that might help them. They are suffering. They are literate. They have an internet connection. They pass through the door another way.

The demographics of the DIY-FMT community are instructive. The most comprehensive survey to date, published by Ekekezie and colleagues in the *American Journal of Gastroenterology* in 2020 [REF:ekekezie2020], sampled 84 self-identified DIY-FMT practitioners through the patient networks of the Peggy Lillis Foundation, the Fecal Transplant Foundation, and the Australian site *The Power of Poop*. The respondents were mostly female (71 percent) and mostly white (92 percent). Most (80 percent) had performed FMT on themselves. Most (92 percent) knew their stool donor personally. Most (87 percent) used internet resources to guide the procedure. The two commonest indications for which DIY-FMT was being used were inflammatory bowel disease (35 percent) and irritable bowel syndrome (29 percent) — both conditions for which, in most jurisdictions, FMT is not offered in regular clinical practice. Adverse events were reported by 12 percent of respondents; 82 percent reported improvement in their condition.

That last figure needs to be looked at carefully, because it is the crux of the DIY-FMT paradox. Eighty-two percent self-reported improvement. That is considerably higher than the remission rate seen in randomised controlled trials of FMT for the same indications (25 to 30 percent in ulcerative colitis, for example). There are at least four possible explanations, and the honest answer probably involves elements of all of them. First, self-report in the absence of blinding is a well-known source of inflated efficacy estimates — patients who have invested substantial effort and risk in a procedure are motivated to report benefit. Second, survey respondents in a voluntary, non-random sample are not representative of all DIY-FMT practitioners; people whose attempts failed or caused harm are under-represented, and people whose attempts succeeded are over-represented. Third, the placebo

effect — which in gastrointestinal disease is notoriously large, sometimes as high as 40 percent in clinical trials — will be operating at full strength in the absence of any blinding or control. And fourth, and most interestingly, the DIY protocols are often substantially *more* intensive than the single-dose protocols used in most randomised trials — weeks of daily enemas from a known, consistent donor, rather than one colonoscopic delivery — and may therefore genuinely produce better engraftment. The multi-dose, intensive DIY protocols resemble, in some respects, the Paramsothy trial’s multi-donor enema schedule, and that was the trial with the best remission rate.

What they actually do

The typical DIY-FMT protocol, as documented in the patient literature on sites like *The Power of Poop*, in the self-help guides circulated through inflammatory-bowel-disease forums, and in the Ekekezie survey, follows a broadly consistent pattern. A donor is recruited — usually a family member or close friend, chosen for apparent good health, absence of antibiotic use in recent months, and willingness. The donor is typically screened informally: a questionnaire, sometimes laboratory screening arranged privately, sometimes nothing at all. A stool sample is produced, usually that morning. It is homogenised in saline, typically in a household blender, then filtered through a coffee filter or a fine muslin. The resulting slurry is administered either as a retention enema via an over-the-counter enema kit, or — less commonly — via a nasogastric-style procedure with a pharmacy-purchased feeding tube. The recipient is instructed to lie down for as long as possible afterwards to retain the material. The procedure is repeated, daily or every few days, often for weeks.

This protocol, read against the modern clinical-grade version, is almost a parody of its professional counterpart. No anaerobic processing: the blender and the coffee filter destroy much of the obligate-anaerobic community. No standardised dose: the recipient does not know how many organisms they are receiving. No donor screening for the pathogens that have, in the clinical setting, killed people: no ESBL, no parasites, often no hepatitis. No quarantine. No post-procedure follow-up. And yet the self-reported outcomes, as we have seen, are not obviously catastrophic — most DIY-FMT practitioners, on survey, report perceiving benefit, and the reported adverse events are usually minor gastrointestinal symptoms rather than the serious infections that clinicians would predict from the absent screening.

Why has the DIY community not produced a visible trail of disease? The most likely answer is that it has — but dispersed through the general infectious-disease case load in ways that do not attribute back to DIY-FMT. A case of ESBL bacteraemia in a DIY-FMT recipient will be diagnosed and reported as ESBL bacteraemia; the DIY-FMT history will be disclosed only if the clinician asks and the patient volunteers it. The literature on adverse events in DIY-FMT is correspondingly thin, not because the events are not happening but because they are going uncounted. The same epistemic problem faces anyone trying to argue, from the absence of a visible harm signal, that DIY-FMT is “safe”. It may not be. We do not know.

What we do know is that, unscreened, it is riskier per procedure than clinical FMT; that it uses donor material that in most cases has never been properly characterised; and that when it fails, the patient has no institutional recourse. A bad outcome from clinical FMT is audited, reported, and contributes to the evidence base. A bad outcome from DIY-FMT usually disappears into the general pattern of illness in the population.

The ethical territory

The existence of DIY-FMT poses a difficult question for the medical profession, and the honest answer is that the profession has not yet settled on a position. The paternalistic response — *this is unsafe, please stop, please trust the system* — is unworkable for two reasons. First, it presupposes that the system is offering something better, which for most non-*C. difficile* indications it is not. Second, it ignores the autonomy of patients who have weighed the risks and made an informed decision about their own bodies, which is a right that medicine has traditionally been slow to challenge when the procedure in question is not FMT. The permissive response — *if you are going to do it anyway, here is how to do it more safely* — runs into regulatory and liability concerns that most institutions are unwilling to confront. The laissez-faire response — *it is not our problem* — is a moral abdication.

The direction that I think the profession is slowly moving, and that I personally find defensible, is something like this: the medical system has an obligation to expand, by responsible research, the range of conditions for which FMT is properly offered within clinical services; to make clinical FMT more accessible for the conditions in which it is indicated (it is currently shockingly hard to obtain in many jurisdictions even for the one indication for which it is the treatment of choice); and, for conditions in which FMT is not formally indicated but a patient is determined to try it anyway, to provide non-judgemental guidance about at least the minimum safety steps — donor screening, above all — that would reduce the risk of catastrophic harm. This is not permission. It is harm reduction, applied to a situation in which the alternative is that patients take risks in ignorance that they could take with some measure of informed caution. Similar harm-reduction frameworks have been accepted, after prolonged resistance, in intravenous drug use, in sexual health, in abortion access. I do not see why FMT should be different.

The reader who is currently considering DIY-FMT for themselves or for a family member — and I know that such readers exist, because I have spoken to them — deserves the following minimum, delivered not as encouragement but as harm reduction. A donor recruited privately should, at an absolute minimum, have a recent negative test for HIV, hepatitis B, hepatitis C, and stool testing for enteric pathogens including *C. difficile*, and should have had no antibiotic exposure in the preceding six months. They should have a normal body-mass-index, no autoimmune or inflammatory bowel disease, no depression, no metabolic syndrome, no history of cancer, and no recent travel to regions with endemic parasitic infection. They should not be sharing intimate contact with the recipient. These tests can be arranged through most primary-care physicians. They do not cost much. They do not guarantee safety. They substantially reduce the risk of the catastrophic outcomes — transmitted bloodstream infection, in particular — that have taken lives in the formal clinical setting. If you are going to do the procedure regardless, do these things. If you are not willing to do these things, reconsider the procedure. That is the honest clinical position, and I deliver it with full awareness that the formal regulatory position of my profession is different.

24.9 From Stool to Standard Product: The End of FMT As We Know It?

The final section of this chapter should describe a future that, in 2026, is arriving faster than most observers predicted. FMT, as I have described it in the preceding sections, is a fundamentally pre-modern therapy in one specific sense: its active ingredient is a biological preparation from a specific human being, whose composition cannot be fully specified, and whose therapeutic effect varies from batch to batch. No other class of drug in modern medicine works like this. Every other biological

therapy — monoclonal antibodies, recombinant proteins, cell therapies — has been progressively standardised, characterised, and manufactured to specification, even when derived from biological sources. FMT is, in this sense, the last large therapy still practised in the nineteenth-century mode: whole crude extract, in bulk, with variable composition, from a variable source.

That is beginning to change. Two products, both derived from the FMT concept but standardised in different ways, received US FDA approval for recurrent *C. difficile* infection in 2022 and 2023 respectively. The first was fecal microbiota, live-jslm (brand name *Rebyota*), a minimally-processed, standardised preparation of pooled donor stool, administered as a single rectal enema, approved in December 2022 and shown in a phase III trial to achieve sustained prevention of *C. difficile* recurrence in approximately 70 percent of treated patients [REF:khanna2025]. The second was SER-109 (brand name *Vowst*), a very different kind of product: not whole-community stool, but a purified preparation of bacterial spores — specifically, the *Firmicutes* spore fraction — prepared from selected healthy donors and administered orally as capsules. SER-109 received FDA approval in April 2023 on the strength of the ECOSPOR III randomised trial, in which it reduced *C. difficile* recurrence to 12 percent over eight weeks, compared with 40 percent in the placebo arm [REF:feuerstadt2022].

These two products are, in a real sense, the post-FMT era in miniature. *Rebyota* preserves the whole-community philosophy of traditional FMT while standardising the preparation, the sourcing, and the quality control. *Vowst* goes much further, reducing the “transplant” to its active spore fraction — the component that is most robust to processing, most clinically defined, and most pharmaceutically tractable — and delivering it as an ordinary oral drug. Both products work. Both were approved. Both are, in their different ways, the first commercial proofs that the concept of microbial community transfer can be turned into a pharmaceutical.

The direction of travel from here is reasonably clear, even if the timescale is not. The next generation of products will likely be *defined consortia*: rationally assembled mixtures of specific, characterised, cultured organisms, designed to reconstitute the key functions of a healthy community without requiring any donor at all. Several such products are in late-phase clinical trials in 2026. Beyond defined consortia are *engineered consortia*, in which specific microbial strains have been modified to carry out particular therapeutic functions — produce a missing metabolite, secrete a defined cytokine, sense a specific environmental cue. And beyond that are the various phage, probiotic, and postbiotic approaches that we have covered in previous chapters. In each case, the underlying therapeutic concept is the one that van Nood’s trial proved in 2013: that modifying the gut microbial community can modify disease. The engineering task is to move from the crude proof-of-concept (a stool suspension from a healthy donor) to the refined implementation (a defined product that does the same thing, better).

If this trajectory continues — and I think it will — faecal microbiota transplantation as a procedure will, over the next decade or two, progressively shrink as a share of the therapeutic landscape. It will remain available for recurrent *C. difficile* where the whole-community approach still outperforms the simpler alternatives (for now), and for research indications where the full complexity of a real human gut is still needed for experimental work. But for most routine clinical applications, the “nuclear option” will be replaced by something less dramatic: a standardised, defined, regulated pharmaceutical product that does what a traditional FMT does, without the donor, without the stool bank, without the enema, and without the public relations problem.

I want to finish this chapter with a return to the metaphor in the title. I said at the beginning that nuclear weapons are blunt, collateral, and irreversible, and that the FMT procedure, on examin-

ation, is none of those things. That was the first mistake of the nuclear metaphor. The second mistake, which I want to name at the end, is that nuclear weapons destroy. FMT, at its best, does not destroy. It restores. A recurrent *C. difficile* patient on their fifth course of vancomycin is inhabiting a gut that has been destroyed, repeatedly, by the antibiotics. The FMT does not add to that destruction. It undoes it. The new community is not a replacement for what was there; it is, in most cases, something very close to what was there before the sequence of illness and treatment took it apart.

A better metaphor, which several of my colleagues have been using and which I am coming around to, is that FMT is a *transplant* in the full medical sense: the transfer of a living biological system from one person to another, with the same careful attention to matching, screening, rejection, and long-term follow-up that we apply to kidneys and hearts and corneas. The sixty-trillion organisms that constitute the microbiome are, after all, an organ by every meaningful criterion — a biomass of a kilogram or more, with discrete functions, tissue-level organisation, a developmental programme, and a role in whole-body physiology. If we can transplant kidneys, we can transplant microbiomes. The surprise of the last fifteen years is not that the procedure works. It is that the organ in question — the one we used to flush away without a second thought — turned out to be an organ at all.

Key Chapter References

[REF:eiseman1958] Eiseman B, Silen W, Bascom GS, Kauvar AJ. Fecal enema as an adjunct in the treatment of pseudomembranous enterocolitis. *Surgery*. 1958;44(5):854-859.

[REF:gough2011] Gough E, Shaikh H, Manges AR. Systematic review of intestinal microbiota transplantation (fecal bacteriotherapy) for recurrent *Clostridium difficile* infection. *Clin Infect Dis*. 2011;53(10):994-1002. doi:10.1093/cid/cir632

[REF:vannood2013] van Nood E, Vrieze A, Nieuwdorp M, et al. Duodenal infusion of donor feces for recurrent *Clostridium difficile*. *N Engl J Med*. 2013;368(5):407-415. doi:10.1056/NEJMoa1205037

[REF:vrieze2012] Vrieze A, Van Nood E, Holleman F, et al. Transfer of intestinal microbiota from lean donors increases insulin sensitivity in individuals with metabolic syndrome. *Gastroenterology*. 2012;143(4):913-916.e7. doi:10.1053/j.gastro.2012.06.031

[REF:alang2015] Alang N, Kelly CR. Weight gain after fecal microbiota transplantation. *Open Forum Infect Dis*. 2015;2(1):ofv004. doi:10.1093/ofid/ofv004

[REF:paramsothy2017] Paramsothy S, Kamm MA, Kaakoush NO, et al. Multidonor intensive faecal microbiota transplantation for active ulcerative colitis: a randomised placebo-controlled trial. *Lancet*. 2017;389(10075):1218-1228. doi:10.1016/S0140-6736(17)30182-4

[REF:cammarota2017] Cammarota G, Ianiro G, Tilg H, et al. European consensus conference on faecal microbiota transplantation in clinical practice. *Gut*. 2017;66(4):569-580. doi:10.1136/gutjnl-2016-313017

[REF:kao2017] Kao D, Roach B, Silva M, et al. Effect of oral capsule- vs colonoscopy-delivered fecal microbiota transplantation on recurrent *Clostridium difficile* infection: a randomized clinical trial. *JAMA*. 2017;318(20):1985-1993. doi:10.1001/jama.2017.17077

[REF:kang2017] Kang D-W, Adams JB, Gregory AC, et al. Microbiota Transfer Therapy alters gut ecosystem and improves gastrointestinal and autism symptoms: an open-label study. *Microbiome*.

2017;5(1):10. doi:10.1186/s40168-016-0225-7

[REF:kang2019] Kang D-W, Adams JB, Coleman DM, et al. Long-term benefit of Microbiota Transfer Therapy on autism symptoms and gut microbiota. *Sci Rep.* 2019;9(1):5821. doi:10.1038/s41598-019-42183-0

[REF:defilipp2019] DeFilipp Z, Bloom PP, Torres Soto M, et al. Drug-resistant *E. coli* bacteremia transmitted by fecal microbiota transplant. *N Engl J Med.* 2019;381(21):2043-2050. doi:10.1056/NEJMoA1910437

[REF:ekekezie2020] Ekekezie C, Perler BK, Wexler A, Duff C, Lillis CJ, Kelly CR. Understanding the scope of do-it-yourself fecal microbiota transplant. *Am J Gastroenterol.* 2020;115(4):603-607. doi:10.14309/ajg.0000000000000499

[REF:ianiro2020covid] Ianiro G, Mullish BH, Kelly CR, et al. Reorganisation of faecal microbiota transplant services during the COVID-19 pandemic. *Gut.* 2020;69(9):1555-1563. doi:10.1136/gutjnl-2020-321829

[REF:ianiro2020] Ianiro G, Bibbò S, Porcari S, et al. Fecal microbiota transplantation for recurrent *C. difficile* infection in patients with inflammatory bowel disease: experience of a large-volume European FMT center. *Gut Microbes.* 2021;13(1):1994832. doi:10.1080/19490976.2021.1994832

[REF:davar2021] Davar D, Dzutsev AK, McCulloch JA, et al. Fecal microbiota transplant overcomes resistance to anti-PD-1 therapy in melanoma patients. *Science.* 2021;371(6529):595-602. doi:10.1126/science.abf3363

[REF:degroot2020] de Groot P, Nikolic T, Pellegrini S, et al. Faecal microbiota transplantation halts progression of human new-onset type 1 diabetes in a randomised controlled trial. *Gut.* 2021;70(1):92-105. doi:10.1136/gutjnl-2020-322630

[REF:feuerstadt2022] Feuerstadt P, Louie TJ, Lashner B, et al. SER-109, an oral microbiome therapy for recurrent *Clostridioides difficile* infection. *N Engl J Med.* 2022;386(3):220-229. doi:10.1056/NEJMoA2106516

[REF:khanna2025] Khanna S, Yoho D, Van Handel D, et al. Prevention of recurrent *C. difficile* infection by fecal microbiota, live-jslm (REBYOTA) administered via colonoscopy: 6-month data from the single-arm phase IIIb CDI-SCOPE trial. *Ther Adv Gastroenterol.* 2025;18:17562848251396744. doi:10.1177/17562848251396744

Chapter 25: Precision Microbiome Engineering — From Nuclear Option to Scalpel

25.1 The Opposite of a Transplant

If Chapter 24 was about the nuclear option, this chapter is about the scalpel. Everything I described in the previous chapter — stool banks, enemas, capsules of pooled donor material, the whole infrastructure that sits between a healthy donor's colon and a sick recipient's — relies, at bottom, on a principle of crude substitution. You take one person's microbial community, more or less whole, and you put it into another person, hoping the new community out-competes the old. It works, for *Clostridioides difficile*, with a reliability that still astonishes me every time I see the numbers. But even its strongest advocates acknowledge that it is a blunt instrument. You cannot, in a faecal transplant, specify which organisms take up residence. You cannot restrict the handover to the species you actually want to deliver. You cannot exclude passengers — antibiotic-resistance genes, latent viruses, metabolic traits — that the recipient did not sign up for. You are, in every clinical sense, accepting a package deal.

Precision microbiome engineering is the opposite ambition. It is the attempt to act on the microbial community at the level of a single species, a single strain, sometimes a single gene within a single strain. Instead of transferring forty thousand bacterial genomes in a bottle, you deliver one carefully chosen organism that does one defined thing. Instead of suppressing a hundred species with an antibiotic so that the one you want to be rid of goes with them, you target that one species specifically, leaving the rest of the community standing. Instead of hoping that a donor's community contains enough of the right metabolic machinery, you engineer the machinery itself and then install it in a chassis organism that knows how to live in the gut.

The distance between these two visions, in 2026, is still substantial. FMT works for a handful of indications; precision engineering works, in humans, for almost none yet. But the trajectory matters. In the fifteen years since the van Nood trial made FMT respectable, precision microbiome engineering has gone from a cartoon on a conference slide to a pipeline of products in human clinical trials, two of them at Phase 3, several of them closer to first-in-human than to last-in-mouse. The field has its first cautionary tales and its first disappointments. It has the beginnings of a regulatory framework. It has a name: *synthetic biology applied to the microbiome*, or, in the language most investors prefer, *live biotherapeutic products*. And it has the characteristic feature of a technology that is about to matter — a lot of smart people working on it, and a lot of money betting that they will be right.

This chapter is the counterpoint to the FMT chapter in the same way that a scalpel is the counterpoint to a blender. I will try, in what follows, to describe what is actually possible in 2026 (which is more than most readers will expect); what is genuinely close but not yet here (which is where most of the excitement lives); what the regulatory apparatus looks like, and why it matters as much as the science; what the specific risks of precision engineering are, including some that are quite different from the risks of FMT; and where the next decade is likely to take us. I will try throughout to distinguish proof of concept from proof of benefit, and both of those from hype. The gap between them is where the reader's caution should live.

25.2 Four Kinds of Precision

Before describing specific products, it helps to lay out the four broad strategies that make up the field. The terminology is not yet fully standardised, and different authors use different taxonomies; the one I find clearest divides precision microbiome engineering into four levels, distinguished by how heavily the organism or the community has been modified.

At the first and least-modified level are the **defined consortia** — rationally chosen mixtures of cultured, naturally occurring bacterial strains, selected from healthy human donors or from culture collections, blended at specified ratios to reproduce the key functions of a healthy gut community. No genetic engineering. No synthesis. Just a curated community, manufactured to specification. These are the direct successors of FMT. *Vowst* (SER-109), which I mentioned at the end of the FMT chapter, is the first-generation example — a mixture of *Firmicutes* spores rather than a true defined consortium, but operating on the same philosophy. The second-generation examples now in late-phase trials, such as VE303 for recurrent *C. difficile* [REF:louie2023], are true defined consortia: specified numbers of specified strains, each one identified, deposited in a culture collection, and manufactured under pharmaceutical conditions.

At the second level are the **single-strain live biotherapeutic products** — specific organisms, delivered on their own, selected because they carry out some particular function the recipient's gut community is lacking. *Akkermansia muciniphila*, the mucus-degrading organism whose abundance correlates with metabolic health, is the prototype. The organism can be grown in culture, pasteurised (interestingly, the pasteurised form appears to retain much of the metabolic benefit), and administered as a capsule. Early randomised trials in overweight, insulin-resistant adults have shown modest improvements in metabolic markers [REF:depommier2019]. The organism is naturally occurring; no engineering is involved. The precision comes from the choice of the single strain, and from the standardisation of the product.

At the third level are the **genetically modified live biotherapeutics** — organisms that have been engineered to perform functions they do not perform in nature. The best-studied chassis is *Escherichia coli* Nissle 1917, a strain isolated by the German physician Alfred Nissle from the stool of a First World War soldier who had, uniquely among his unit, not developed dysentery during the Balkans campaign. Nissle 1917 is a safe, well-characterised commensal that has been used as a probiotic for a century. In the last decade it has become the *E. coli* K-12 of the microbiome — the default workhorse into which you install the genes you want expressed in the human gut. Synlogic, a Massachusetts company spun out of MIT in the mid-2010s, built an entire pipeline on Nissle 1917 chassis engineered to consume metabolites that accumulate to toxic levels in specific inherited metabolic diseases. I will describe their progress, and their setbacks, in a moment.

At the fourth and most aggressive level are the **sequence-specific antimicrobials** — therapies that use CRISPR-Cas systems, or other sequence-guided mechanisms, to kill specific bacterial strains within a community while leaving the rest untouched. These are in many ways the most “surgical” of all the precision approaches, because they act not on the recipient's metabolism but on the genome of the organism they are targeting. Two companies, Locus Biosciences in North Carolina and Eligo Bioscience in Paris, have been the pioneers; both have delivery platforms based on engineered bacteriophages carrying CRISPR payloads, and both have published or reported clinical-trial data in the last two years. We will come back to them.

These four levels are not entirely separate. A defined consortium can contain an engineered organism; a CRISPR-based antimicrobial can be delivered alongside a live biotherapeutic that moves into the space the target vacates. The real pipeline combines them. But the taxonomy helps sort the story, and I will use it.

25.3 What Already Works: The Defined Consortia

Of the four levels, the one with the most established clinical foothold in 2026 is the defined consortium. This is not a coincidence. Defined consortia are, in a regulatory sense, the easiest to approve: the organisms are naturally occurring, the manufacturing processes are adapted from food-industry fermentation practice, and the conceptual link to FMT — a treatment already shown to work for the target indication — provides a ready-made efficacy rationale. In the post-FMT regulatory environment that the FDA, the European Medicines Agency, and their counterparts have been constructing since around 2019, defined consortia have been the first class of microbiome therapy to cross the finish line.

VE303 and the second wave in recurrent *C. difficile*

I described *Vowst* and *Rebyota* at the end of the FMT chapter as the first commercial proofs that community-transfer could be pharmaceuticalised. Both were derived from donor stool; neither was synthetic in the purist sense. The product that, if it crosses its remaining regulatory hurdles, will be the first genuinely defined consortium to reach market is VE303, from the Boston company Vedanta Biosciences. VE303 is a mixture of eight specific *Clostridium* strains, originally isolated from the stool of a single healthy adult donor, cultured in pure form, banked as master cell lines, and manufactured as a mixture of freeze-dried spores in a capsule. Every strain in VE303 has a genome sequence in a public database. Every batch is verified by sequencing to contain exactly those eight strains in specified ratios. The product is, in this sense, a pharmaceutical in the way that *Rebyota* — derived from a pool of real human donors whose exact composition cannot be fully specified — is not quite yet.

The pivotal Phase 2 trial of VE303 for recurrent *C. difficile* infection, reported in 2023, randomised seventy-nine patients to either high-dose VE303, low-dose VE303, or placebo, each given as a daily oral capsule regimen for fourteen days after standard antibiotic treatment [REF:louie2023]. The recurrence rate over eight weeks was 13.8 percent in the high-dose arm, 37 percent in the low-dose arm, and 45.5 percent in the placebo arm. The confidence interval for the high-dose comparison just excluded zero, but the pattern — dose-dependent, ecologically plausible, and of a magnitude comparable to *Vowst* and *Rebyota* — was considered convincing enough that the product entered a Phase 3 programme. That programme was still running at the time of writing; an approval in 2026 or 2027 is plausible but not certain.

What makes VE303 more than just a third entry in the *C. difficile* market is its conceptual status. For the first time, a community-based therapy is being approved on the basis of *exactly which organisms are in it*, not on the basis of *which donor it came from*. This is the first post-donor microbiome therapy. If VE303 is approved, the model of microbiome engineering that it represents — identify the active components from a natural community, culture them separately, manufacture them as a defined mixture — becomes the default template for the next generation. A handful of similar products, targeting inflammatory bowel disease, hepatic encephalopathy, and cancer immunotherapy augmentation, are behind VE303 in the pipeline.

The inflammatory bowel disease consortium front

The FMT literature in ulcerative colitis that I reviewed in Chapter 24 showed a real but modest signal — remission rates around 25 to 30 percent against 5 to 10 percent with placebo, with the best results from the most intensive multi-donor pooled protocols. The defined-consortium answer to that literature is to ask: what are the specific organisms in the pool that are actually driving the benefit, and can we reconstruct those and omit the rest?

Several groups have taken a stab at this question. The 4D Pharma consortium MRx0518, built around a single *Enterococcus gallinarum* isolate; the Finch Therapeutics product CP101, a capsulised whole-stool preparation that was closer to FMT than to true consortium engineering; the VE202 consortium, sixteen strains designed to induce regulatory T cells. Not all of them have prospered. Finch Therapeutics filed for bankruptcy in 2023 after failing to secure funding for its Phase 3 programme, a reminder that the commercial environment for microbiome therapeutics remains unforgiving. MRx0518 and VE202 have both had trials that have produced mixed data, and both are undergoing protocol revisions at the time of writing.

The broader lesson from this front is worth stating plainly. The technology of building defined consortia works. The difficulty is that the diseases being targeted — ulcerative colitis, Crohn's disease, irritable bowel syndrome — are more biologically complex than recurrent *C. difficile*, in ways that no amount of microbiome manipulation by itself is likely to fully resolve. A defined consortium that restores community diversity may modify these diseases; it is unlikely, on current evidence, to cure them. Realism about what the technology can do is going to matter, in the next five years, as much as the technology itself.

25.4 The Engineered Commensals: Synlogic and the Bacteria That Eat Poisons

The second major precision-engineering strategy — genetic modification of a safe commensal to carry out a defined therapeutic function — has had a more complicated history. The most visible pipeline has been Synlogic's, and the Synlogic story is worth telling in some detail because it illustrates, in a single company's arc, what can be done with this technology, what remains hard, and what happens when a genuinely elegant scientific idea meets the blunt reality of clinical trial design.

The Synlogic strategy, in a sentence, is this: take a safe, well-characterised commensal (*E. coli* Nissle 1917), engineer it to metabolise a toxic compound that accumulates in a specific inborn error of metabolism, deliver it as an oral capsule, and use the engineered bacterium as a living “metabolic sink” that converts the toxin into something harmless in the patient's gut, before the toxin can be absorbed systemically. The elegance of the idea is that it turns a cell into a drug: not in the metaphorical sense in which all cell therapies are drugs, but in the literal sense that the cell has been given an enzymatic activity it did not previously have, and that activity is the therapeutic mechanism.

SYNB1020 and the hyperammonaemia setback

The first Synlogic product to reach clinical trials was SYNB1020, a Nissle 1917 strain engineered to consume ammonia and convert it to arginine. The target indication was urea-cycle disorders and hepatic encephalopathy — conditions in which ammonia accumulates in the blood to neurotoxic levels because the normal liver pathway for ammonia disposal is either genetically defective

(urea-cycle disorders) or overwhelmed by liver failure (encephalopathy). A Phase 1 trial in healthy volunteers demonstrated that the engineered strain transited the gut, expressed the engineered pathway, and altered nitrogen metabolism in a direction consistent with the intended mechanism [REF:kurtz2019]. The science worked.

The Phase 1b/2a trial in patients with cirrhosis and hyperammonaemia, however, did not meet its efficacy endpoint and was terminated in 2019. The reasons are instructive. The engineered organism performed its metabolic function as designed; the problem was that it did not perform enough of it, for long enough, in the right place, to move the clinical endpoint. The gut is a harsh and crowded ecosystem. An introduced organism, even a well-adapted one like Nissle 1917, does not colonise stably in most adults; it transits through over days, doing what engineering it has while it is present, and is then out-competed by the resident community. The “dose” of the therapeutic, in the sense of how much ammonia-consuming enzyme is actually acting in the relevant anatomical compartment for the relevant amount of time, turned out to be lower than mouse models had suggested. This is a recurring lesson in the engineered-commensal space: the mouse gut and the human gut are not equivalent bioreactors, and the extrapolation from one to the other is still being learned.

SYNB1618 and phenylketonuria

The second Synlogic product, SYNB1618, was engineered to degrade phenylalanine — the amino acid that accumulates to neurotoxic levels in phenylketonuria (PKU), the inherited metabolic disease familiar from the heel-prick test given to every newborn in developed countries. PKU patients live, for their entire lives, on a diet that excludes most protein-containing foods, a regime that is particularly punishing for children and adolescents. An oral therapeutic that could degrade dietary phenylalanine in the gut before it could be absorbed — effectively decoupling diet from blood phenylalanine levels — would be transformative.

SYNB1618, and its successor SYNB1934, made it through Phase 2 trials with a signal of benefit in blood phenylalanine reduction. But the Phase 3 trial, known as Synpheny-3, was terminated in 2023 on the recommendation of the data monitoring committee, after an interim analysis suggested that the study was unlikely to meet its primary endpoint at the planned sample size [REF:synlogic2023]. The termination was not because the product did not work — it did work, in the sense of reducing blood phenylalanine modestly — but because the effect was too small relative to placebo to be clinically compelling, and because during the trial’s life an alternative product (pegvaliase, an injectable phenylalanine-degrading enzyme) had moved onto the market and made the bar for approval higher. Synlogic announced, shortly afterwards, that it was winding down its engineered-bacterium programmes and re-focusing on other technologies. The company’s share price, which had exceeded ten dollars at its peak, fell below a dollar.

The Synlogic story is not, however, a failure of the underlying science. The trials demonstrated that engineered *E. coli* Nissle 1917 could be administered safely to humans over long periods; that the engineered metabolic pathways functioned in the human gut; that the organisms did not colonise persistently (a safety feature, if a therapeutic limitation); and that the immunological, metabolic, and microbial-ecological consequences of the intervention were manageable. The lesson is not that engineered commensals do not work, but that the specific combination of delivery, residence time, and enzymatic capacity required to move a systemic clinical endpoint is a harder engineering problem than the founders initially estimated. Subsequent groups are attacking it with different chassis organisms (spore-forming *Bacillus* species, which survive the stomach more robustly;

gut-adapted *Bacteroides* species, which colonise persistently if they can be made safe), different delivery schedules, and different indications. The Synlogic programme may have closed, but the engineered-commensal field is not, in 2026, demonstrably closing with it.

The wider engineered-commensal pipeline

Other companies and academic groups are pursuing the same general strategy for different targets. Novome Biotechnologies engineered a *Bacteroides* strain to degrade oxalate, the kidney-stone-forming metabolite that accumulates in hyperoxaluria; their programme reached Phase 1b but has since paused for funding reasons. ZBiotics, which operates in the consumer space rather than as a prescription pharmaceutical, has an engineered *Bacillus* product marketed to degrade acetaldehyde after alcohol consumption (a hangover prevention product, with modest evidence). The academic group of Timothy Lu at MIT, which originally span Synlogic out, has continued to publish on new chassis organisms and new genetic-circuit designs that offer more stable expression and more sophisticated behaviour — including “sense-and-respond” circuits in which the engineered organism only expresses its therapeutic payload when it detects a specific signal in its environment. None of these has yet reached the market. Several are plausibly five to ten years from doing so.

The direction of travel in this subfield is clear enough. The first generation of engineered commensals was built on the question “can we do this at all?” and the answer was, roughly, yes. The second generation is being built on the question “can we do this well enough to move a clinical endpoint?” and the answer depends on the target. For diseases where a relatively small metabolic perturbation is enough to matter — certain inborn errors of metabolism, gut-luminal toxins, locally-acting immune modulators — the technology is plausibly close. For diseases that require sustained, systemic, large-magnitude metabolic shifts, the engineering problem remains hard.

25.5 The CRISPR Antimicrobials: Killing One Species at a Time

The fourth precision-engineering strategy is, in my reading, the most technologically radical, and the one with the most transformative near-term potential if the engineering problems can be solved. It is also, probably not coincidentally, the one that most clearly illustrates why precision engineering is genuinely different from antibiotics, probiotics, or FMT.

The observation on which the strategy rests is that CRISPR-Cas systems — the bacterial immune systems we have spent the last fifteen years repurposing for gene editing in human and plant cells — are, in their native context, sequence-specific antimicrobials. A bacterium encoding a CRISPR-Cas system that has captured a sequence from a phage it previously encountered will, on re-encountering that phage, cut the phage DNA to pieces. The specificity is set by a short guide RNA that matches twenty nucleotides of the target sequence; the cutting is done by the Cas nuclease. Repurpose the system so that the guide RNA matches a unique, essential sequence in the genome of a target pathogen rather than a phage, and the Cas nuclease will cut the pathogen’s own chromosome, causing the bacterium to lyse. Deliver the CRISPR-Cas system to the target pathogen using a bacteriophage that naturally infects that pathogen’s species, and you have a sequence-specific antibiotic: it kills only organisms whose genomes match the guide RNA, and spares everything else in the community.

This is, in the language of medicine, an astonishing prospect. Every other antimicrobial we have — every antibiotic, every antifungal, every antiviral, every phage — is either broad-spectrum (and therefore collateral) or narrow-spectrum in a species-at-best way. A CRISPR antimicrobial can be,

in principle, strain-specific: it can kill the pathogenic *Escherichia coli* that produces an extended-spectrum beta-lactamase while leaving the commensal *Escherichia coli* that does not. It can eliminate an antibiotic-resistance gene from a community without killing any of the organisms carrying it (a variant strategy, using Cas systems that cut plasmids rather than chromosomes, that is also being pursued). It can target a virulence factor rather than a species. It is, in the deepest sense, a therapy that can be programmed.

Locus Biosciences and LBP-ECO1

Locus Biosciences, founded in North Carolina in 2015 and partly acquired by Swiss pharmaceutical company Roche in 2024, has been the most clinically advanced of the CRISPR-antimicrobial companies. Their lead product, LBP-ECO1, is a cocktail of three engineered bacteriophages, each carrying a CRISPR-Cas3 system with guide RNAs targeting essential and virulence-related sequences in uropathogenic *E. coli* — the most common cause of urinary tract infection. The choice of Cas3 rather than Cas9 is deliberate: Cas3 is a processive nuclease that shreds DNA progressively along the chromosome after the initial cut, making bacterial escape through repair or mutation much less likely than with Cas9.

In 2024, Locus reported the results of its ELIMINATE Phase 2 trial in patients with recurrent urinary tract infections [REF:locus2024]. The trial met its primary safety endpoint — no drug-related serious adverse events — and showed a reduction in *E. coli* burden in the urinary tract over seven days of treatment compared with placebo, with the specificity of the effect (preservation of other urinary and gut commensals) confirmed by sequencing. A Phase 3 programme is now in preparation. If LBP-ECO1 is approved — and in 2026 this looks plausible within the next two to three years — it will be the first CRISPR-based antimicrobial to reach clinical practice. The implications extend well beyond urinary tract infection. Drug-resistant *E. coli* is one of the five bacterial species responsible for the majority of deaths from antimicrobial resistance globally [REF:murray2022]; a therapeutic that can selectively eliminate the pathogenic strains without damaging the commensal community would, at scale, be a substantial addition to the global antimicrobial-resistance toolkit.

Eligo Bioscience and the “eligobiotic” platform

The French company Eligo Bioscience, founded in 2014 by scientists at the Rockefeller University and the Pasteur Institute, has pursued a similar concept with a different delivery strategy. Rather than using intact engineered phages, Eligo uses phage-derived capsids that deliver a CRISPR-Cas9 payload but cannot replicate within their target cells — a design that should, in principle, reduce the risk of the therapy propagating uncontrollably through the community. Their lead programme targets the toxin-producing *Staphylococcus aureus* strains implicated in certain skin and systemic infections, with a particular focus on strains producing Pantone-Valentine leukocidin. A further programme addresses the colibactin-producing strains of *E. coli* that have recently been implicated in colorectal cancer initiation. First-in-human trials for the lead programme were scheduled to begin in late 2025 or early 2026 [REF:eligo2025]. As of this writing, their status is not publicly confirmed.

Eligo also reported, in 2023, the first *in vivo* demonstration of microbiome base-editing in a mammalian gut — using a non-cutting, sequence-specific deaminase delivered via phage capsid to alter a specific DNA position in a target strain of *E. coli* in the intestines of live mice, without killing the organism or any of its neighbours. The result is conceptually important. It demonstrates that the CRISPR-toolbox can be applied, in the crowded and chaotic environment of a real mammalian mi-

crobiome, not just to destroy specific organisms but to *edit* them in place — leaving the organism alive but altering a specific function. What that opens up, in therapeutic terms, is not yet clear. But the demonstration itself is a proof of principle that will, in the next decade, be iterated on by many groups.

The specific risk profile of CRISPR antimicrobials

The precision of CRISPR antimicrobials is their greatest advantage. It is also the source of their characteristic risks, which differ in interesting ways from the risks of conventional antimicrobials.

First, resistance. A sequence-specific antimicrobial can, in principle, be defeated by a single mutation in the target sequence. In practice, the careful design of multi-target cocktails (LBP-ECO1 has three phages each targeting different sequences) and the choice of essential, conserved genes as targets minimise this risk, but it is real and will require monitoring in long-term use. Second, off-target effects. A guide RNA sequence chosen to match a unique site in the target pathogen must be verified, computationally and experimentally, not to match other sites in other organisms in the community. Bioinformatic screening for off-target matches is now routine, but the reference microbiome data on which such screening depends is incomplete, and novel off-target effects in an individual patient's specific microbial community cannot be ruled out in advance. Third, phage-host immunology. The bacteriophage delivery vehicle is itself a foreign biological entity, and repeated administration can induce host antibodies that reduce the effectiveness of subsequent doses — the same problem that AAV-based human gene therapies face, adapted to the microbial context. Fourth, horizontal gene transfer. The most serious theoretical concern is that a CRISPR-bearing bacteriophage might pass its payload to a non-target organism via horizontal transfer, producing unintended effects elsewhere in the community. The engineering strategies in current use (replication-incompetent capsids, narrow-host-range phages, inducible systems) are designed specifically to minimise this risk, but the field is not yet mature enough to be certain that all routes of unintended propagation have been closed.

None of these risks is disqualifying. All of them are specific to precision antimicrobials, and all of them are addressable with engineering attention. The regulatory frameworks are still catching up; the CRISPR antimicrobial is a genuinely novel regulatory object, neither drug in the traditional small-molecule sense nor cell therapy in the gene-therapy sense, and the agencies are feeling their way.

25.6 The Regulatory Landscape: Why Approval Is Slower Than the Science

At this point in the chapter I want to pivot from the science to the regulation, because in a field as new as this one the regulatory trajectory is not a downstream consequence of the science; it is a load-bearing determinant of what can be developed, for what indications, and how quickly. The reader who has followed pharmaceutical development at all is familiar with the FDA/EMA/TGA model for small-molecule drugs: Phase 1 for safety, Phase 2 for preliminary efficacy, Phase 3 for pivotal efficacy, then filing for approval, then post-marketing surveillance. Microbiome therapeutics fit awkwardly into this model, and the regulatory agencies have, over the last decade, been building a parallel framework that inherits some of the conventional drug model and some of the biologics and cell-therapy models.

The FDA's Live Biotherapeutic Product framework

The FDA's principal category for microbiome therapeutics is the **Live Biotherapeutic Product (LBP)**, a term introduced in agency guidance documents around 2016 and refined several times since. An LBP is a biological product that contains live organisms — bacteria, fungi, or, in some readings, phage — that is applicable to the prevention, treatment, or cure of a disease in humans, and that is not a vaccine. LBPs are regulated by the FDA's Center for Biologics Evaluation and Research (CBER), the same centre that handles vaccines and blood products, rather than by the drug centre. The LBP framework imposes a set of manufacturing standards (current Good Manufacturing Practice, adapted for live-organism processes), a set of characterisation requirements (whole-genome sequencing of master cell banks, identity and purity testing of each batch, stability testing to establish shelf-life), and a clinical-trial framework that is broadly similar to that for small-molecule drugs with some additions for the live-organism dimension (assessment of shedding from treated subjects, long-term follow-up for unanticipated ecological effects).

The LBP framework covers defined consortia and single-strain products but does not neatly accommodate either genetically modified organisms (which also come under the gene-therapy framework, adding complexity) or phage-based CRISPR antimicrobials (which sit at the intersection of LBP, gene-therapy, and antimicrobial regulation). In practice, these more complex products require case-by-case regulatory pathways, with considerable bilateral discussion between sponsor and agency about what data are required and what endpoints are acceptable. This case-by-case negotiation adds years to development programmes.

The European and Australian contexts

The European Medicines Agency has a broadly similar framework under the heading of Advanced Therapy Medicinal Products (ATMPs), with extensions to accommodate microbial therapeutics. The UK's Medicines and Healthcare products Regulatory Agency (MHRA), post-Brexit, has been notable for somewhat greater flexibility, positioning itself explicitly as a jurisdiction that will consider novel microbiome therapeutics on expedited pathways where the science warrants it. Australia's Therapeutic Goods Administration has, to this date, regulated FMT under its Biologicals framework and has been slower to build a parallel LBP regime; this has, in practice, resulted in most Australian patients being treated with imported products rather than domestically developed ones. China's National Medical Products Administration approved its first defined-consortium product in 2024 on the basis of Chinese-only clinical-trial data, and has positioned itself to compete with the FDA for LBP approvals in the second half of this decade. Regulatory fragmentation, as with most pharmaceutical categories, means that a single product often has three or four different approval pathways in different markets, with substantial duplication of effort.

The specific regulatory complications of precision engineering

Beyond the general complications of live-organism therapies, precision microbiome engineering creates specific regulatory challenges that the agencies are still working through.

The first is the question of genetic containment. A genetically modified live organism, if it persists or spreads in the environment, is in principle a piece of genetically modified material released without controlled end-point — a category that, for agricultural organisms, is highly regulated. Applied to a therapeutic bacterium in the gut of a treated patient, the same logic generates questions about shedding, environmental release, and possible horizontal gene transfer to the general microbial population. The current regulatory response is to require that engineered strains be auxotrophic

(dependent on a nutrient they cannot synthesise and that is not available in the wild) so that they cannot survive outside the treated gut, and to require tracking studies showing that engineered DNA does not persist in the recipient's stool beyond the treatment period. These requirements are scientifically reasonable but add considerable complexity and cost to the engineering itself.

The second is the question of off-target effects in a living, heterogeneous community. A small-molecule drug can be characterised *in vitro* for its target specificity and its off-target binding; a CRISPR antimicrobial can be characterised for its guide-RNA sequence specificity; but the behaviour of either in a specific patient's actual microbial community cannot be fully predicted. The regulators are, understandably, asking for more and more patient-level community characterisation as part of trial designs, with 16S and shotgun sequencing of stool and other sites at multiple timepoints. This requirement, which did not exist for first-generation FMT approvals, is now standard for precision products. It is generating a large amount of data whose interpretation is not yet settled.

The third is the question of long-term surveillance. The residence time of an introduced organism in the gut can range from days (for most non-colonising commensals) to years (for engineered organisms specifically designed to colonise persistently). The regulators are, quite reasonably, asking for much longer post-marketing surveillance periods for precision-engineered products than for conventional biologics — up to ten years for some engineered-commensal products. This requirement is defensible but commercially discouraging.

The net effect of all these considerations is that precision microbiome therapeutics, while scientifically advancing quickly, are moving through clinical development at a slower pace than either their underlying science or their commercial advocates would prefer. This is probably, in the main, a good thing. The lessons of early gene therapy — the deaths in the 1999 Jesse Gelsinger trial, the leukaemias in the French SCID-X1 trial of 2002 — are what taught the field the value of regulatory conservatism when the therapy is both novel and biologically complex. Microbiome engineering is at a similar stage, and will probably benefit from being made to go slowly.

25.7 The Risks Specific to Precision Engineering

Throughout the preceding sections I have noted individual risk profiles for each strategy. Let me now draw them together, because the risk landscape of precision microbiome engineering is meaningfully different from, rather than simply smaller than, the risk landscape of FMT.

The risks of FMT, as the previous chapter described, are dominated by the transfer of unintended biological content from a human donor to a human recipient — pathogens, resistance genes, possibly metabolic or behavioural traits. The safety fix is better donor screening, and the residual risk is small but real. Precision engineering sidesteps these risks almost entirely: a manufactured product derived from defined strains or engineered organisms has, by design, no undefined passengers. The DeFilipp ESBL death that reshaped FMT safety could not, in principle, happen with VE303.

But precision engineering introduces its own set of risks, some of which are genuinely new to medicine.

First, **unintended ecological cascade**. A carefully chosen targeted intervention in a complex ecosystem can have consequences that the intervention's designers did not anticipate, because the ecosystem contains feedback loops and compensatory mechanisms that are not captured in any

current model. Kill a target pathogen with a CRISPR antimicrobial, and another organism may expand into the niche the pathogen vacated. Deliver an engineered commensal that produces a specific metabolite, and the metabolite may alter the behaviour of organisms elsewhere in the gut that were not anatomically or functionally considered in the engineering brief. Mouse models and computational simulations help, but the ecosystem in a specific patient at a specific time is always partly unknowable in advance. Adverse events in this category are, by construction, hard to predict. In the clinical trials to date they have not been a dominant signal, but the follow-up periods are too short to be confident.

Second, **horizontal gene transfer**. Bacteria exchange genetic material through plasmid conjugation, phage transduction, and natural transformation. An engineered gene delivered to one organism in the gut has, in principle, the possibility of being taken up by another organism and persisting, at low frequency, in the broader microbial population. The engineering strategies in use (chromosomal integration rather than plasmid carriage, elimination of mobile genetic elements from the chassis, use of degradation-targeting sequences) minimise this risk, but do not eliminate it. The consequence, if it occurred, would depend on what gene was transferred and to what organism. It could be benign; it could, in unlucky circumstances, create a novel resistance or virulence profile in a new host.

Third, **unintended activation of the host immune system**. Live bacteria, even probiotic ones, can trigger immune responses. Engineered organisms carrying novel protein sequences (the fluorescent marker proteins sometimes used for tracking, or the therapeutic enzymes themselves) have an additional source of potential immunogenicity. The Synlogic trials did not report serious immune-related adverse events, but the therapy was given for limited periods; a product intended for long-term or lifelong administration would have different stakes.

Fourth, **off-target genome editing in the patient's own cells**. This is, in most cases, a theoretical concern rather than a demonstrated risk. CRISPR-Cas systems delivered by phage to bacterial targets are operating inside bacteria, not in the patient's own cells, and the mechanism of delivery (phage infection) is biologically specific to bacteria. But the underlying concern — that CRISPR machinery in the gut might somehow reach a human cell and edit a human gene — is one that the regulators are asking sponsors to address explicitly, with tracking data showing that the editing machinery does not escape its intended compartment.

Fifth, and in some respects most subtly, **manufacturing consistency**. Live organisms are biological products, and biological products have lot-to-lot variability. A batch of VE303 that happens to have 2 percent lower levels of a specific strain than the specification calls for is not obviously defective, but it is, in ecological terms, a different product from the one characterised in the pivotal trial. The quality-control systems that the LBP framework requires are designed to minimise this variability, but they do not eliminate it. What happens when individual patients receive out-of-specification batches, and how such variability would be detected in routine clinical practice, are questions still being worked out.

Sixth, **cost and access**. This is not a biological risk, but it is a risk to the population-level benefit of the technology. The pharmaceutical manufacturing of live engineered products is expensive — more expensive than the manufacturing of most small-molecule drugs, and probably expensive on an ongoing basis rather than converging to low cost as patents expire. The market prices of *Vowst* and *Rebyota* in the US, at the time of writing, are in the tens of thousands of dollars per course. Whether precision microbiome therapies become generally accessible, or whether they join the list of modern therapeutics that are technically available but practically restricted to wealthy patients

in wealthy countries, is a question that will be settled by health-system economics rather than by science.

25.8 What Is Plausibly Near, and What Is Not

I want to end with my own best reading of what is genuinely close, what is genuinely possible but further out, and what belongs more to the next generation than to this one. I have tried throughout this chapter to be sober; here I will be more so.

Near-term (2026–2030). Additional defined consortia will be approved for recurrent *C. difficile* infection, including VE303 or a successor. The first CRISPR-based antimicrobial, probably LBP-EC01 or a close competitor, will be approved, most likely for a narrow indication such as recurrent urinary tract infection caused by a specific *E. coli*. The first defined-consortium product for ulcerative colitis may reach approval, with modest efficacy comparable to the FMT trials. One or two engineered commensals for inborn errors of metabolism will reach late-phase trials, though approval will depend on effect sizes that have, so far, been smaller than hoped.

Medium-term (2030–2035). The move from whole-community FMT to defined consortia will largely complete for the major indications, leaving true FMT as a procedure mainly of historical and salvage importance. CRISPR antimicrobials will have been deployed, possibly at scale, for at least one antimicrobial-resistance problem (likely *E. coli* bloodstream infection or hospital-acquired *Klebsiella pneumoniae*). Engineered commensals will probably have achieved approval for at least one inborn error of metabolism. Early trials of “live sensor” organisms — engineered strains that detect specific biomarkers in the gut and either report them externally or respond with a therapeutic payload — will be underway. Base-editing of resident organisms in vivo, currently a laboratory proof of concept, may reach first-in-human trials.

Longer-term (2035–2045). The field’s central ambition — programmable microbial communities in which composition and function can be both measured and modified with precision comparable to current genetic medicine — will be approaching clinical reality for a narrow range of indications. Whole-community design, rather than single-strain modification, may become tractable as the underlying ecology is better understood. Some of the grander speculations — microbial communities engineered for the prevention of cardiovascular disease, for the modulation of ageing, for the alteration of behavioural phenotypes — will be either experimentally demonstrated or quietly abandoned, and it is not yet possible to say which of them will be which.

I have three cautionary notes to add to this optimistic arc. The first is that the gap between proof of concept and proof of benefit, in this field, is repeatedly turning out to be larger than its advocates expected. Synlogic’s arc is the cautionary tale here; the engineering worked, but the clinical benefit was too modest, too many times. The second is that precision microbiome engineering sits in a broader biomedical environment in which alternative technologies — monoclonal antibodies, small-molecule drugs, RNA therapeutics — are continuously improving, and the bar a microbiome therapy has to clear is rising. A microbiome therapy that would have been transformative if developed in 2015 may be merely useful if developed in 2030, because its competition has improved in the meantime. The third is that the underlying biology is still being learned. Every few years, a mechanism is uncovered that changes what we thought we knew — recent examples include the role of bacteriophage communities in shaping bacterial population structure in the gut, the recognition that small RNA signals travel between organisms in ways that were not previously modelled,

and the emerging appreciation of biofilm structure in gut ecology. The microbiome is not a bag of bacteria. It is an ecosystem with layered structure, and every layer has surprises left in it.

25.9 Scalpel, Not Nuclear: A Final Framing

I opened this chapter with the contrast between nuclear option and scalpel, and I want to return to it at the end, because I think the metaphor captures something important about the choice that this field offers the reader.

FMT is, in a real sense, a brute-force solution to a problem we do not yet fully understand. It works because a complex adaptive system — a human gut community — has enough redundancy, enough internal cross-talk, and enough resilience that when you replace the whole system at once you can, under the right conditions, restore health. What you cannot do, with a nuclear option, is direct the restoration. You cannot say: please give me back the bile-acid-metabolising species that I lost to antibiotics, but leave everything else as it was. You take the whole donor community or none of it.

Precision engineering is the ambition to direct the restoration. To say, in effect: these specific organisms, performing these specific functions, in these specific amounts. It is, when it works, a fundamentally more sophisticated kind of medicine than FMT. It matches the way we think about modern pharmacology — active ingredients, dose-response, selectivity, characterisation. It also, when it does not work, fails in characteristic modern-pharmacology ways: the mechanism is right, the delivery is inadequate; the target is correct, the effect size is small; the drug does what was promised, but the disease has too many other causes for that one action to be enough.

The right reader to end this chapter with is the one who will, in a decade or two, have a choice between FMT and a precision product for some indication not yet approved. That reader, if they ask whether the precision product is *better*, is asking a question that the answer will depend on. Better at being characterised, standardised, regulated, manufactured at scale, and shipped without a donor: certainly. Better at accessing the full compensatory biology of a real human gut community, with its tens of thousands of genes and its millennia of coevolutionary fine-tuning: not necessarily, and in some cases probably not. A scalpel is more precise than a blender, but a blender covers more ground. The case for precision engineering is not that it makes FMT obsolete. It is that it complements FMT; that it takes the parts of the microbiome intervention space where the precision matters (specific metabolic functions, specific resistance elimination, specific single-organism delivery) and carries them forward into a pharmaceutical age, while leaving the parts where whole-community complexity matters — and they may turn out to be more of them than we currently think — to the more ancient technology of the transplant.

The next decade of microbiome therapeutics will, if I am reading the field correctly, see not the replacement of FMT by precision engineering but the bifurcation of the field into two parallel tracks. Precision therapies for indications where we understand the biology well enough to design a specific intervention; whole-community therapies for indications where we do not, and where the surest way to restore function is still to restore community. Both tracks will mature. Neither will make the other irrelevant. And the patient in front of me in ten years' time will, I hope, be offered whichever of the two is right for their specific disease, their specific community, and their specific place in the long and still only-beginning story of how humans learned to live, deliberately, with the organisms that live with them.

Key Chapter References

- [REF:depommier2019] Depommier C, Everard A, Druart C, et al. Supplementation with *Akkermansia muciniphila* in overweight and obese human volunteers: a proof-of-concept exploratory study. *Nat Med*. 2019;25(7):1096-1103. doi:10.1038/s41591-019-0495-2
- [REF:kurtz2019] Kurtz CB, Millet YA, Puurunen MK, et al. An engineered *E. coli* Nissle improves hyperammonemia and survival in mice and shows dose-dependent exposure in healthy humans. *Sci Transl Med*. 2019;11(475):eaau7975. doi:10.1126/scitranslmed.aau7975
- [REF:louie2023] Louie T, Golan Y, Khanna S, et al. VE303, a defined bacterial consortium, for prevention of recurrent *Clostridioides difficile* infection: a randomized clinical trial. *JAMA*. 2023;329(16):1356-1366. doi:10.1001/jama.2023.4314
- [REF:synlogic2023] Synlogic Inc. Synlogic announces discontinuation of Synpheny-3 Phase 3 clinical trial of labafenogene marselecobac (SYNB1934) for the treatment of phenylketonuria. Press release, 18 August 2023.
- [REF:locus2024] Kim P, Sanborn M, Nixon J, et al. A randomized, double-blind, placebo-controlled Phase 1b/2a study of LBP-EC01, a CRISPR-Cas3-enhanced bacteriophage cocktail, for the suppression of *E. coli* in the urinary tract of patients with recurrent urinary tract infections (ELIMINATE trial). [Company-reported conference data and clinical trial registration NCT05488340.]
- [REF:eligo2025] Eligo Bioscience. Pipeline and investor update: EBO05 programme advancing towards first-in-human studies. Company release, 2025.
- [REF:murray2022] Murray CJL, Ikuta KS, Sharara F, et al. Global burden of bacterial antimicrobial resistance in 2019: a systematic analysis. *Lancet*. 2022;399(10325):629-655. doi:10.1016/S0140-6736(21)02724-0
- [REF:lagier2016] Lagier J-C, Khelaifia S, Alou MT, et al. Culture of previously uncultured members of the human gut microbiota by culturomics. *Nat Microbiol*. 2016;1:16203. doi:10.1038/nmicrobiol.2016.203
- [REF:bikard2014] Bikard D, Euler CW, Jiang W, et al. Exploiting CRISPR-Cas nucleases to produce sequence-specific antimicrobials. *Nat Biotechnol*. 2014;32(11):1146-1150. doi:10.1038/nbt.3043
- [REF:citorik2014] Citorik RJ, Mimee M, Lu TK. Sequence-specific antimicrobials using efficiently delivered RNA-guided nucleases. *Nat Biotechnol*. 2014;32(11):1141-1145. doi:10.1038/nbt.3011
- [REF:neil2021] Neil K, Allard N, Roy P, et al. High-efficiency delivery of CRISPR-Cas9 by engineered probiotics enables precise microbiome editing. *Mol Syst Biol*. 2021;17(10):e10335. doi:10.15252/msb.202110335
- [REF:pdaia2023] FDA. Early clinical trials with live biotherapeutic products: chemistry, manufacturing, and control information — guidance for industry. U.S. Food and Drug Administration, Center for Biologics Evaluation and Research. 2016 (updated 2023).
- [REF:hanson2021] Hanson BM, Weinstock GM. The importance of the microbiome in epidemiologic research. *Ann Epidemiol*. 2021;35:75-79. doi:10.1016/j.annepidem.2019.05.007

Chapter 26: Inflammatory and Autoimmune Disease

26.0 The Twentieth-Century Curve

There is a graph that has shaped a generation of immunological thinking. On the vertical axis is the prevalence of chronic inflammatory and autoimmune diseases — inflammatory bowel disease, type 1 diabetes, multiple sclerosis, rheumatoid arthritis, coeliac disease, asthma, allergic rhinitis, atopic dermatitis. On the horizontal axis is time, from roughly 1950 to the present. Every one of those curves is rising, most of them by a factor of two to four, in industrialised countries where the rise cannot be attributed to better diagnosis or longer life expectancy. In some conditions — paediatric Crohn’s disease, paediatric type 1 diabetes — the incidence in children has more than doubled in the space of a single generation.

These diseases are genetically grounded. They cluster in families, they are linked to specific human leukocyte antigen (HLA) alleles and to hundreds of smaller-effect genetic variants identified by genome-wide association studies. But genetics alone cannot explain the epidemiology. The human genome has not changed meaningfully in two generations. Something in the environment is pushing susceptible people over a threshold into clinical disease, and the faster the rise, the more the “environment” must be doing relative to the genes.

The older candidates for the environmental trigger — diet, infection, vitamin D, pollution — remain on the table, and this chapter will not dismiss any of them. But the candidate that has most reshaped the research agenda in the last two decades is the microbiome. The rising curves of autoimmune disease line up, disconcertingly well, with a separate and partly independent set of observations: that diversity of the human gut microbiota has been falling in industrialised populations, that antibiotic exposure in early life has become near-universal, that Caesarean delivery rates have climbed, that the first six months of infant feeding have moved away from the ecological pattern under which our species evolved. If the microbiome shapes the developing immune system — and the preceding chapters have argued, in detail, that it does — then a systematic shift in the community we carry could be one of the things tipping genetically susceptible people into disease.

That hypothesis is what this chapter examines. It is not yet proven, and in some of the diseases we will discuss it is less proven than in others. But it is no longer a fringe idea. The leading journals of microbiology, immunology, and gastroenterology have, over the last five years, published synthesis after synthesis arguing that inflammatory and autoimmune diseases are best understood as failures of host–microbe negotiation rather than as failures of the immune system in isolation [REF:iliev2025][REF:kim2023]. The five diseases I have chosen as case studies — inflammatory bowel disease, type 1 diabetes, multiple sclerosis, rheumatoid arthritis, and coeliac disease — each tell the story from a different angle. Taken together, they sketch the outline of a common mechanism: an ecological shift in the gut community, a weakening of the epithelial barrier that normally contains it, a chronic low-grade leak of microbial products into tissues where they do not belong, and an immune system that was trained by one community and is now being asked to police another.

That common mechanism, if it holds, has two properties that matter for the rest of this chapter.

The first is that it implies some of these diseases might be preventable or modifiable by ecological interventions — dietary, probiotic, or more radical. The second is that it does not promise any such intervention is yet available. A hypothesis about pathogenesis is not a treatment, and the gap between “we think we understand the mechanism” and “we can fix it in a patient” is wider, in this field, than the popular press tends to acknowledge. The chapter will try to hold both of these truths in view at once: the genuine explanatory power of the microbiome model, and the still-modest evidence that it can be translated into clinical benefit for the individual patient.

One final note before we begin. Of all the chapters in this book, this is the one in which the marketing overlay is most dangerous. “Leaky gut,” “autoimmune protocol,” “gut reset,” and a hundred variants are sold to patients with real and serious diseases, by practitioners who borrow the language of the peer-reviewed literature while detaching it from the evidence that language is supposed to rest on. I will use some of those phrases in this chapter, because they point to real phenomena, but I will try to draw the line, every time, between the scientifically defensible claim and the commercial version. The reader who finishes this chapter will not be equipped to diagnose themselves or prescribe their own treatment. They will, I hope, be equipped to read the next magazine article about “fixing your microbiome to cure your autoimmune disease” with the scepticism it deserves.

26.1 Inflammatory Bowel Disease: The Paradigm Condition

If the microbiome hypothesis of autoimmune disease has a home ground, it is inflammatory bowel disease.

IBD is an umbrella term for two related but distinct conditions. *Crohn’s disease* can affect any part of the gastrointestinal tract from the mouth to the anus, but most commonly the terminal ileum and the colon, producing transmural inflammation — inflammation that goes all the way through the bowel wall — with a characteristic pattern of “skip lesions” interspersed with healthy tissue. *Ulcerative colitis* is confined to the colon, involves the mucosa and submucosa only, and produces a continuous inflammatory front that begins in the rectum and extends proximally. Both diseases are chronic, relapsing, and, in their severe forms, life-altering. They require long-term immunosuppression. They carry increased risks of surgery, of colorectal cancer, and of a range of systemic complications from arthritis to uveitis to primary sclerosing cholangitis of the bile ducts.

IBD was the first major disease in which the microbiome was seriously implicated, and it remains the one in which the evidence for microbial involvement is strongest. There are several reasons why IBD became the paradigm.

The first is that the disease occurs in the gut — the place where the microbiome and the immune system meet, and where a disturbed community can most plausibly translate into clinical disease. You do not need a complicated story about gut-to-brain or gut-to-joint signalling to explain why dysbiosis might cause colitis. The microbes are right there, separated from the host immune system by a single layer of epithelial cells and a fragile, replaceable mucus blanket. When that barrier fails, the consequences are immediate and local.

The second is that the genetics of IBD pointed at the microbiome from the beginning. The first major susceptibility gene identified for Crohn’s disease, *NOD2* (originally *CARD15*), encodes an intracellular sensor for muramyl dipeptide — a fragment of the bacterial cell wall. People with loss-of-function variants in *NOD2* cannot properly sense this bacterial signal, and they are three to thirty times more likely to develop Crohn’s disease, depending on the variant [REF:hugot2001].

The subsequent decade of IBD genetics identified more than two hundred additional loci, and the genes are heavily enriched for functions related to innate immunity, epithelial barrier integrity, autophagy of intracellular bacteria, and the handling of microbial signals. The genetic architecture of IBD is, in large part, the genetic architecture of the host–microbe interface. A disease defined this way almost has to involve the microbiome.

The third reason is practical. The organ in question is one we can sample. Stool is easy to collect; colonoscopic biopsies, while invasive, are part of standard IBD care and are routinely available for research. The IBD microbiome has therefore been sequenced at a scale that no other autoimmune disease can match, and the findings have been replicated across dozens of independent cohorts on three continents.

What the data show

The bacterial signature of IBD, when you cut through the noise of different methodologies and different populations, is strikingly consistent. It has three main components.

The first is *reduced diversity*. IBD patients have fewer bacterial species in their guts than healthy controls, and the species they carry are less evenly distributed. The Faith phylogenetic diversity metric and the Shannon diversity index — two ways of capturing how “rich” a microbial community is — are significantly lower in both Crohn’s disease and ulcerative colitis, and the deficit persists even in clinical remission, when the bowel has healed.

The second is *depletion of specific beneficial organisms*. The most consistently reduced organism in both forms of IBD is *Faecalibacterium prausnitzii* — the butyrate-producing keystone species we met in Chapter 23 as one of the most desirable residents of the healthy colon. In IBD, *F. prausnitzii* falls from its normal abundance of five to fifteen per cent of the community to, in active disease, often less than one per cent. Other butyrate producers — *Roseburia intestinalis*, members of the *Clostridium* clusters IV and XIVa — show similar depletion. Butyrate, as we discussed in Chapter 20, is the main fuel for colonocytes and a potent inducer of regulatory T cells. A gut that cannot produce butyrate is a gut whose lining is undernourished and whose immune tone has lost one of its main pacifiers.

The third is *expansion of pro-inflammatory organisms*. Two have attracted particular attention. The first is *adherent-invasive* Escherichia coli (AIEC), a pathotype of *E. coli* that is able to adhere to and invade intestinal epithelial cells and, unusually, to survive and replicate inside macrophages [REF:darfeuille2004]. AIEC is overrepresented in the ileal mucosa of Crohn’s disease patients, and the *NOD2* loss-of-function variants that predispose to Crohn’s are precisely the variants that impair the intracellular handling of this kind of organism. You can see the logic: a host whose immune system cannot contain AIEC gives AIEC the freedom to expand and cause chronic inflammation. The second is *Ruminococcus gnavus*, an anaerobic gut commensal that, in IBD, expands dramatically and produces a glucorhamnan polysaccharide that triggers inflammatory cytokine release from dendritic cells [REF:henke2019]. *R. gnavus* blooms are particularly common during IBD flares, and recent work has begun to suggest that the strains that bloom are genetically distinct from the strains carried by healthy people — a pathogenic ecotype hiding within what was thought to be a benign species.

These findings are robust. They replicate across studies and populations. If you sequence the stool of a hundred IBD patients and a hundred controls, you will reliably find lower diversity, lower *F. prausnitzii*, and higher pro-inflammatory organisms in the patient group. This is about as solid as

observational microbiome science gets — and the pattern is summarised cleanly in the 2023 review by Haneishi and colleagues, which is a useful entry point for any reader wanting a more technical overview than the one I am providing here [REF:haneishi2023].

The Gevers cohort: catching the microbiome at the moment of disease onset

But observational data from established disease cannot tell you whether the microbiome differences cause the disease or merely result from it. Inflammation changes the gut environment — pH, oxygen tension, bile acid flux, mucosal integrity — in ways that favour some organisms and punish others. A community that looks dysbiotic in an IBD patient might be the *consequence* of the inflammation, not its cause. To distinguish the two, you need to study the microbiome *before or at the moment of disease onset*, ideally before any treatment has further confused the picture.

That is what Dirk Gevers and colleagues at the Broad Institute did in the largest paediatric Crohn's disease microbiome study to date, published in *Cell Host & Microbe* in 2014 and still one of the foundational papers of IBD microbiome research [REF:gevers2014]. The RISK cohort enrolled children at the time of their diagnostic colonoscopy, before they had received any immunosuppressive treatment. Samples were collected from multiple sites — the ileum, the rectum, and stool — which matters because, as we saw in the Zmora probiotic study, stool does not always tell the same story as mucosal biopsy. The team sequenced the microbiomes of 447 new-onset Crohn's patients and 221 controls.

The findings replicated and sharpened what was already known. There was a specific “axis” of dysbiosis that tracked tightly with disease status: increased *Enterobacteriaceae*, *Pasteurellaceae*, *Veillonellaceae*, and *Fusobacteriaceae*, and decreased *Erysipelotrichales*, *Bacteroidales*, and *Clostridiales*. The signal was stronger in the ileum and the rectum than in the stool — a finding that, like Zmora's, reminds us that stool is a proxy, not a substitute. And crucially, the signature was already present at diagnosis, before any drug had altered the picture. Whatever else the IBD microbiome is, it is not *solely* a downstream consequence of treatment.

One further observation from the Gevers paper has shaped subsequent research more than the primary findings did. Children who had been given antibiotics in the weeks before their diagnostic colonoscopy had a *more* dysbiotic picture than those who had not — the antibiotics, often prescribed for a presumed infection before the correct diagnosis was reached, had amplified the underlying dysbiosis rather than treating it. This is a cautionary note about the interpretation of the IBD microbiome in clinical practice, and it is also a first hint that antibiotic exposure might be more than a confounder: it might be a contributor.

The antibiotic connection

That hint has since hardened. Multiple large population-based studies have found that antibiotic exposure in the first year of life is an independent risk factor for the subsequent development of paediatric IBD, with odds ratios in the range of 1.3 to 2.0 — modest individually, but meaningful across a population in which early-life antibiotic use is the rule rather than the exception. A 2020 Swedish registry study of nearly a million children found that each additional course of antibiotics in the first two years of life increased IBD risk by around fifteen per cent, and that the effect was strongest for agents with the broadest spectrum. The biological interpretation, consistent with the themes of Chapter 18, is that antibiotics disrupt the founding community during the critical developmental window, leaving a gut ecology that is more vulnerable to the combination of genetic susceptibility and dietary or other environmental triggers that Crohn's disease requires to manifest.

This is not a reason to avoid necessary antibiotics in children. It is a reason to be thoughtful about unnecessary ones — and, as I argued in Chapter 18, the proportion of paediatric antibiotic prescriptions that are clinically unnecessary is uncomfortably high.

The cross-kingdom reframe

Until recently, the IBD microbiome story was told almost entirely in bacterial terms. That is beginning to change. A 2025 review in *Nature Reviews Microbiology* by Iliya Iliev, Ashwin Ananthakrishnan and Chun-Jun Guo represents, in my view, the most important synthesis of the IBD microbiome field of the past decade [REF:iliev2025]. The Iliev review argues, and I think successfully, that the IBD microbiome is not a bacterial community with fungi and viruses as minor contaminants. It is a *cross-kingdom* community — bacteria, fungi, and bacteriophages interacting with each other and with the host immune system — and its pathological shifts are cross-kingdom shifts. The bacterial dysbiosis is real, but it is part of a larger ecological disturbance that also involves the fungal community (which we will examine in the next section) and the viral community (which Chapter 17 introduced).

The review's central claim is that IBD therapy has focused on the bacterial piece because that is the piece we understand, but that the other pieces — the mycobiome, the virome — may be where the unexploited therapeutic opportunities lie. Some of those opportunities are already in early trials: phage cocktails targeting AIEC, antifungal approaches for *Candida*-dominant mycobiome shifts, bacterial consortia designed to restore *F. prausnitzii* and other butyrate producers. Whether any of these will become standard care in the next decade is uncertain. That they are being tried at all is a measure of how far the field has come.

A word about primary sclerosing cholangitis

Before leaving IBD, I want to say something about a condition that sits at its edge and that the chapter outline does not explicitly cover: primary sclerosing cholangitis (PSC). PSC is a chronic inflammatory disease of the bile ducts, progressively scarring and narrowing them, ultimately producing liver cirrhosis and the need for transplantation. About seventy per cent of PSC patients also have IBD — usually ulcerative colitis — and the association is so strong and so specific that the two are best thought of as facets of a single broader inflammatory syndrome.

The microbiome of PSC is distinct from that of IBD without PSC. A 2024 study in the *Journal of Crohn's and Colitis* compared gut microbiomes and bile acid profiles in 54 IBD-PSC patients and 62 IBD-only patients, and found that the IBD-PSC group had reduced microbial gene richness, distinct taxonomic shifts (decreased *Blautia obeum*, increased *Veillonella atypica*, *Veillonella dispar*, and *Clostridium scindens*), and increased microbial capacity for *secondary bile acid metabolism*, with corresponding shifts in serum bile acid profiles [REF:leibovitzh2024]. The mechanistic story, emerging from this and other work, is that gut bacteria modify the bile acids the liver has produced, and that abnormal microbial bile acid metabolism may drive the cholangiopathy of PSC. It is a striking example of a principle that matters for the whole chapter: the microbiome does not cause disease only in the organ it lives in. Through metabolic transformations of host molecules, it can cause disease anywhere the blood carries those molecules.

I mention PSC here because it is the clearest single example in this chapter of microbial metabolism driving disease at a distance. The liver is not a microbial organ. It has no residents of its own. But the bile it secretes is reshaped by the gut community, and when that reshaping goes wrong, the bile ducts are where the damage appears. The gut-liver axis, like the gut-brain axis we met in Chapter

13, is a real and clinically consequential communication channel, and the PSC story is where it most clearly tips into disease.

26.2 The Fungal Dimension: The ASCA Mystery, Thirty Years On

In Chapter 10 I promised that we would come back to one of the stranger stories in modern gastroenterology: the *Anti-Saccharomyces cerevisiae antibody* — ASCA — which was discovered in the early 1990s as a serological marker of Crohn's disease and which, in the intervening thirty years, has quietly changed the way we think about the gut's fungal residents. The payoff is overdue. This is where the cross-kingdom reframe of §26.1 becomes concrete, and where a question the reader has been carrying since Chapter 10 receives something approaching an answer.

The mystery

Saccharomyces cerevisiae is baker's yeast. It is the organism that makes bread rise and beer ferment. It is not a pathogen. It has never been seriously proposed as a cause of inflammatory bowel disease. And yet, when researchers in Lille and Rouen in the early 1990s were searching for serological markers that could distinguish Crohn's disease from ulcerative colitis — a clinically useful goal, because the two conditions require different long-term management — they found that a large fraction of Crohn's patients carried circulating antibodies that bound to the cell wall of *S. cerevisiae*. About sixty per cent of Crohn's patients were ASCA-positive. Only about ten per cent of ulcerative colitis patients and almost no healthy controls were.

The finding was useful. ASCA rapidly entered clinical practice as part of the diagnostic workup in IBD, helping to sort patients with ambiguous presentations into one diagnostic bucket or the other. But the underlying question — *why do Crohn's patients make antibodies against baker's yeast?* — went unanswered for a long time. Baker's yeast is in bread, and every healthy person on earth eats bread (or equivalents), and yet only Crohn's patients mount a significant antibody response against it. What was the immune system actually responding to?

The answer, pieced together over the next two decades and beautifully summarised in a 2023 *Autoimmunity Reviews* paper by Boualem Sendid and colleagues from the original Lille group [REF:sendid2023], is molecular mimicry. The epitope that ASCA binds on *S. cerevisiae* is an oligomannose — a short chain of mannose sugars — displayed on the yeast's cell wall glycoprotein coat. And the same oligomannose, or something close enough to cross-react with the same antibody, is displayed on the cell wall of *Candida albicans*.

Candida albicans is not baker's yeast. It is, as we saw in Chapter 10, a routine resident of the human gut mycobiome. In most people it is a benign commensal that lives quietly among the bacteria, participating in the cross-kingdom community without causing trouble. But in the gut of a Crohn's disease patient, *Candida* tends to overgrow. The fungal community shifts. The barrier becomes leakier. And the immune system begins to see, in abundance, a fungal antigen that it would normally encounter only in small doses.

The ASCA antibody, on this reading, is not really an antibody against baker's yeast. It is an antibody against *Candida* — or against the shared oligomannose epitope that *Candida* and *S. cerevisiae* happen to display — and the only reason we found it by screening against baker's yeast is that baker's yeast was the antigen sitting in our reagent freezers when the original studies were designed. It is

a fingerprint left in the immune system's memory by chronic, low-grade exposure to a gut fungus that has grown past its niche.

I find this a remarkably satisfying piece of scientific history. It explains, in a stroke, the long-standing puzzle of why an antibody against a dietary yeast would mark a disease that the dietary yeast clearly does not cause. It links the IBD serology to the IBD mycobiome. And it provides a gentler, more ecological reading of the immune system's behaviour: not an aberrant response to an irrelevant antigen, but a reasonable response to an antigen that has become, for ecological reasons, too abundant to ignore.

The fingerprint-reader analogy

If that seems abstract, consider an analogy. Your immune system is, among other things, a fingerprint reader. Over a lifetime of exposure to the organisms living on and in you, it builds up templates — B-cell repertoires, T-cell clones, antibody specificities — that match the shapes it has seen most often. The templates are built from the most immunogenic features of those organisms, and they are kept ready to deploy if those features appear in threatening quantities. The system is, for the most part, tuned to recognise the specific residents of that specific body.

When a Crohn's patient's gut begins to host more *Candida* than it should, their immune system builds a template to match the *Candida* cell wall. Then, when you put them in front of a reagent freezer full of baker's yeast — a different yeast, but one that happens to share a cell-wall motif with *Candida* — the template lights up. We see it light up, and we call the result ASCA. We name the antibody after the organism that showed the signal, not after the organism that trained the response.

This kind of cross-reactivity is not rare. It is how most vaccines work (exposure to a deactivated pathogen trains a response that recognises the live one) and it is the basis of many autoimmune reactions (an immune response to a microbe happens to cross-react with a structurally similar host antigen). In this chapter we will see it again, in the discussion of coeliac disease and in the putative mechanisms of rheumatoid arthritis. Molecular mimicry is not a rare exotic mechanism. It is a routine feature of how the adaptive immune system works, with occasional pathological consequences when the cross-reacting antigen is a self-antigen or a component of a tissue that was never meant to be attacked.

The Candida story consolidates

The mycobiome evidence in IBD has, over the past decade, consolidated around the ASCA narrative. A systematic review published in *Inflammatory Bowel Diseases* in 2025 by Johannes Fuchtbauer and colleagues pooled twenty-seven case-control mycobiome studies comprising 1,406 IBD patients [REF:fuchtbauer2025]; an earlier but complementary narrative review by Hsu and colleagues in the same journal in 2023 reaches a largely concordant conclusion and is useful for readers who want a deeper dive into the mechanistic and therapeutic implications [REF:hsu2023]. The findings were striking for their consistency, at the higher taxonomic levels, and striking for their *inconsistency* at the species and strain levels — a combination that turns out to be characteristic of mycobiome research.

At the higher level: *Candida* species are reliably increased in both Crohn's disease and ulcerative colitis, *Saccharomyces* species are reliably decreased, and *Malassezia* is specifically increased in Crohn's disease, sometimes dramatically. The bacterial picture of reduced diversity and reduced *Faecalibacterium* is mirrored, in other words, by a fungal picture of reduced *Saccharomyces* and

increased *Candida*, with both sides of the ledger pointing the same direction: away from the commensal configuration and toward the inflammatory one.

At the lower level — which specific *Candida* species, in which proportions, in which patients — the evidence is messier. Some studies find *C. albicans* dominant; others find *C. tropicalis* or *C. parapsilosis*. The methodological heterogeneity is part of the story. Mycobiome studies use different primers, different sequencing platforms, different bioinformatics pipelines, and different thresholds for calling a species “present.” A large part of the apparent disagreement between studies reflects measurement differences rather than biological differences. The review authors are commendably honest about this, and their call for methodological standardisation is one of the strongest sub-texts of the paper.

The *Malassezia* finding deserves a note of its own. Until quite recently, *Malassezia* was thought of as a strictly skin-surface fungus — the organism responsible for dandruff and for some forms of seborrhoeic dermatitis, and a normal resident of the sebum-rich regions of the skin that we met in Chapter 6. The discovery that *Malassezia* is significantly overrepresented in the inflamed mucosa of Crohn’s disease patients was, at first, a surprise. Further work has confirmed that specific *Malassezia* species — notably *M. restricta* — can colonise the gut, engage the innate immune receptor CARD9 (another Crohn’s-susceptibility gene), and drive pro-inflammatory cytokine production. The fungal kingdom, like the bacterial kingdom, is more portable than we thought.

From association to intervention

The therapeutic implication of the mycobiome story is that antifungal interventions, or strategies that restore *Saccharomyces* to the community, might have a role in IBD management. The evidence for this is, as of 2026, preliminary. *Saccharomyces boulardii* — the probiotic yeast we met in Chapter 23 — has been tested in several small IBD trials, with modest and inconsistent results. Antifungal drugs have not been systematically evaluated in IBD, because the conventional view, until recently, was that fungi were bystanders rather than drivers. That view is changing, and over the next few years we can expect more rigorous trials.

A cautionary note: the cross-kingdom framework does not mean that everyone with Crohn’s disease should be prescribed an antifungal. *Candida albicans* is a normal gut resident in most people, and a perspective piece by Shao, Haslam, Bennett and Way in *Trends in Immunology* argues convincingly that the organism should be understood not as a pathogen that is occasionally tolerated but as a *symbiont whose interactions with the host immune system are, in health, mutually beneficial* [REF:shao2022]. It is its overgrowth and the loss of its bacterial competitors that matter, not its mere presence. Simply killing *Candida* without addressing the underlying ecological disturbance is, predictably, a short-term strategy. The broader lesson of Chapter 23 — that ecological context, not species identity, determines outcomes — applies with full force here. “Antifungal therapy for Crohn’s” is a marketing phrase in search of an evidence base. When that evidence base arrives, it will almost certainly be strain-specific, patient-specific, and context-dependent, in the same way that probiotic evidence has turned out to be.

What the ASCA story offers the reader, above the technical details, is a worked example of what cross-kingdom thinking looks like when it pays off. For thirty years we measured the antibody without understanding what it was telling us. Once we allowed ourselves to think about the gut mycobiome as an actual community rather than an afterthought, the antibody’s meaning clarified, and with it a piece of the IBD puzzle that had been missing. The chapter you are reading, and the book it sits inside, are an argument for applying that kind of thinking to more of the autoimmune

diseases on our twentieth-century curve. The next three sections try to do exactly that.

26.3 Type 1 Diabetes and the “Leaky Gut” Hypothesis

Type 1 diabetes (T1D) is a disease in which the immune system destroys the insulin-producing beta cells of the pancreas. It usually appears in childhood or adolescence, it is lifelong, and it absolutely requires exogenous insulin to manage — without it, the patient dies within days or weeks. It has a strong but incomplete genetic component, concentrated in the same HLA region that predisposes to coeliac disease and to several other autoimmune conditions. And, like the other diseases on our twentieth-century curve, its incidence is rising far faster than the genome can account for. In much of Northern Europe, T1D incidence in children has roughly doubled in the past thirty years.

The speed of that increase has driven a search for environmental triggers, and the microbiome has emerged, over the past fifteen years, as the leading hypothesis. The mechanism proposed is elegant, if still incompletely proven, and it centres on a concept whose name — “leaky gut” — has been both a legitimate scientific term and one of the most abused phrases in contemporary health marketing.

What “leaky gut” actually means

The gut lining is a single layer of epithelial cells sealed together by specialised protein junctions — most importantly the *tight junctions* that run around the apical edges of each cell like a rubber gasket around a jar lid. The tight junctions allow water and small molecules to pass between cells (the paracellular route) while preventing larger molecules, bacteria, and microbial products from crossing. Together with the mucus layer, the epithelial cells, and the immune cells beneath, they form the intestinal barrier that keeps the microbial community in the lumen and the host tissue on the other side.

The barrier is not perfectly impermeable. It is a regulated barrier, opening and closing to let selected molecules through — nutrients, immune signals — under the control of a range of cell-biological mechanisms. One of those mechanisms is controlled by a small protein called *zonulin*, first identified by Alessio Fasano and colleagues at the University of Maryland in the late 1990s and early 2000s. Fasano’s work established that zonulin is the main physiological modulator of human tight-junction permeability, that its secretion can be triggered by certain bacterial exposures and by gliadin (the gluten protein, which we will return to in §26.5), and that increased zonulin levels correlate with measurable increases in intestinal permeability [REF:fasano2012].

When the barrier is intact and the zonulin system is functioning normally, the tight junctions regulate themselves, the barrier holds, and the bacterial world and the host world remain, mostly, separate. When the barrier is compromised — through dysbiosis, through inflammation, through specific dietary triggers in susceptible individuals, or through infection — microbial products that normally stay in the lumen begin to cross into the lamina propria, into local lymphatics, and from there into the systemic circulation. The most studied of these translocating products is *lipopolysaccharide* (LPS), the outer-membrane component of Gram-negative bacteria that potently activates the innate immune system through toll-like receptor 4. Low-grade systemic LPS exposure is one of the clearest biomarkers of the leaky-gut state, and it is the mechanism most often invoked to link gut dysbiosis to inflammation in tissues that do not themselves carry a microbiome.

This is the scientifically defensible version of the leaky-gut hypothesis: measurable increases in paracellular permeability, measurable translocation of microbial products, and

measurable downstream immune activation. It is a real phenomenon, and the machinery has been worked out in enough detail to make it a credible player in autoimmune disease [REF:divincenzo2023][REF:mu2017].

The marketing version — which holds that “leaky gut” is a specific named syndrome that causes essentially every modern illness and that can be cured by the proprietary supplement on offer — is a different creature, and the reader should not confuse the two. The phenomenon is real. The claim that it is *the* cause of your fatigue, your allergies, your brain fog, your autoimmune disease, and your weight gain, and that Dr So-and-So’s twenty-eight-day protocol can fix it, is a claim that cannot be supported by the current evidence, and in many cases is actively contradicted by it. A cautious 2023 review surveying the supplements most often marketed for “treating leaky gut” — glutamine, zinc, collagen peptides, various polyphenols and prebiotic fibres — found that the human clinical evidence for any of them is, at best, preliminary; most of the supporting data come from cell-culture and rodent studies whose translation to human disease is unproven [REF:aleman2023]. I will use the phrase “leaky gut” in this section because the biological phenomenon deserves a name and this is the name it has been given. But I am using it in the Fasano sense, not in the Instagram sense, and I will distinguish the two whenever the risk of confusion is high.

TEDDY: the natural history of the pre-diabetic microbiome

The most important body of evidence on the microbiome and type 1 diabetes comes from a single extraordinary study. The Environmental Determinants of Diabetes in the Young — TEDDY — is a multi-centre prospective cohort in which thousands of children at elevated genetic risk for T1D have been followed from birth, with stool samples collected monthly and with detailed clinical phenotyping across six US and European sites. It is, to my knowledge, the most comprehensive longitudinal study of the human microbiome in relation to any disease, and the papers that have emerged from it are foundational for this section.

In October 2018, the TEDDY group published two companion papers in *Nature*, one by Stewart and colleagues on the normal development of the gut microbiome in early childhood [REF:stewart2018teddy], and one by Vatanen and colleagues on the microbiome in children who went on to develop islet autoimmunity or overt T1D [REF:vatanen2018]. Between them, the two papers characterised more than 10,000 stool samples from 783 children, using both 16S ribosomal RNA gene sequencing and shotgun metagenomics. They are, in a sense, the reference dataset against which all subsequent T1D microbiome work is compared.

The Stewart paper established the baseline. The healthy infant gut microbiome passes through three distinct phases: a *developmental* phase from roughly three to fourteen months, during which the community is low-diversity and dominated by *Bifidobacterium*; a *transitional* phase from fifteen to thirty months, during which diversity rises rapidly and the community shifts toward a more adult-like configuration; and a *stable* phase from thirty-one to forty-six months in which the community approaches its adult form. The single most important factor shaping the developmental phase was breast milk exposure, which supported high *Bifidobacterium* abundance (particularly *B. breve* and *B. bifidum*) and slowed the maturation trajectory toward the adult configuration. Mode of birth, geography, siblings, and pet exposure all contributed, but breastfeeding dominated. This is, essentially, the biology we have been describing throughout the book, captured at unprecedented resolution in a large cohort.

The Vatanen paper then asked the interesting question. In a nested case-control analysis, were the children who went on to develop islet autoimmunity or T1D microbiologically distinguishable from

those who did not — before seroconversion, before any clinical sign of disease? The finding was subtle but consistent. The microbiomes of children who remained disease-free contained more genes involved in fermentation and short-chain fatty acid biosynthesis than the microbiomes of children who went on to develop autoimmunity. The taxonomic signal was diffuse — no single species reliably distinguished future cases from future controls — but the *functional* signal, at the level of what the community was metabolically capable of doing, was coherent across sites. The gut communities of children who would later develop T1D were, on average, less well equipped to produce the short-chain fatty acids that, as we saw in Chapter 20, fuel colonocytes and induce regulatory T cells.

The authors stopped short of claiming causation. The study is observational, the effects are modest, and the mechanistic links between SCFA production and pancreatic beta-cell protection, while biologically plausible (and supported by mouse work in NOD diabetic models), are not yet proven in humans. But the direction of the evidence is clear: a functionally impoverished gut community, particularly one low in SCFA producers, precedes T1D by months to years, and it does so in a way that is consistent with the broader autoimmune-microbiome model.

The subsequent decade of TEDDY and TEDDY-adjacent work has added further layers. Children who developed T1D were more likely to have been exposed to antibiotics in the first year of life. They were more likely to have had a dysbiotic pattern at the time of introduction of solid foods. Their microbiomes showed elevated *Bacteroides* and reduced *Bifidobacterium* prior to seroconversion, and their blood biomarkers showed evidence of increased intestinal permeability (higher zonulin, higher LPS-binding protein) compared to matched controls [REF:luppi2024][REF:asantebaadu2025]. The picture that emerges — consistent across several independent cohorts — is of an altered gut ecosystem producing less butyrate and propionate, a weakened epithelial barrier, and low-grade chronic translocation of microbial products, which taken together bias the developing immune system toward loss of tolerance to self-antigens, including the islet antigens that are the eventual targets in T1D.

The hygiene hypothesis, updated

Readers of a certain age will recognise the shape of this argument. It is a late-model descendant of the “hygiene hypothesis,” first articulated by David Strachan in 1989 as an explanation for the rising prevalence of hay fever and allergic disease in Western populations. Strachan’s original claim was that early-life exposure to infections, and to the sibling interactions that deliver them, trained the immune system away from allergic responses. In the decades since, the hypothesis has been reformulated, re-reformulated, and is now usually called the “microbial deprivation hypothesis” or the “old friends hypothesis” to avoid the implication that cleanliness itself is the problem.

The updated version is closer to what this chapter is arguing. It is not that infections themselves protect against autoimmune disease. It is that the broad, continuous, low-level microbial exposure under which our immune system evolved — the community of bacteria, fungi, phages, parasites, and dust-dwelling organisms that our ancestors lived intimately among for hundreds of thousands of years — calibrated the immune system in ways that modern hygienic, antibiotic-saturated, urban life cannot replicate. We have not lost our immune system. We have lost the microbial training partners that shaped it, and we are beginning to see what an untrained immune system does when exposed to the usual range of human antigens. In susceptible individuals, it attacks them.

T1D fits this framing well. It is rare in populations still living under traditional microbial conditions. It rises rapidly in populations that adopt Western sanitation, Western obstetric practices, Western

antibiotic prescribing, and Western infant feeding. The microbiome data from TEDDY and related cohorts describe, in biological detail, what the “old friends” framing predicts at the epidemiological level: a developing gut community that has been thinned and functionally impoverished, and that is no longer able to induce the regulatory T cell responses required for tolerance to self.

The limits of the model

I have to offer the same caution here that applies to every section of this chapter. The association is reproducible. The mechanism is biologically plausible. The early-life timing is consistent with the “critical window” biology we discussed in Chapter 3. But association and mechanism do not, on their own, make a therapy.

No intervention trial to date has shown that modifying the infant microbiome prevents type 1 diabetes. Probiotics in at-risk infants have been tested in small studies with inconsistent results, and none has met the evidentiary standard that would be required to change clinical practice. The Finnish ABIS cohort has tested extensively hydrolysed versus standard infant formulas for T1D prevention, with broadly negative results. The hypothesis is good. The *intervention* — at the population level or at the individual level — has not yet been identified.

There is, as I warned at the beginning of this chapter, substantial commercial pressure around this space. Products marketed as preventing or reversing T1D via “gut health” are, at the time of writing, scientifically untested in the relevant age group and unsupported by any rigorous evidence. A parent whose child has just been diagnosed with T1D is in an emotionally extraordinary state, and the temptation to try something — anything — that might offer hope is entirely human. But the microbiome story, however compelling at the level of biology, does not yet offer an intervention that changes the clinical course of the disease. The current role of this research is to deepen our understanding of pathogenesis, not to guide treatment. That distinction is easily lost in translation from the journal to the magazine article to the supplement bottle.

What the TEDDY data do allow is a kind of realistic optimism. If the pre-disease microbiome signature can be identified early enough, and if a safe intervention can be designed that restores the missing functional capacity (SCFA production, barrier function, regulatory T cell induction), then prevention becomes at least conceivable. Conceivable is not the same as feasible, and feasible is not the same as proven. But conceivable is further than we were a decade ago, and the Vatanen paper is a major part of the reason.

26.4 Multiple Sclerosis and Rheumatoid Arthritis: Autoimmunity at a Distance

Inflammatory bowel disease is autoimmunity in the gut. Type 1 diabetes is autoimmunity in the pancreas, an organ adjacent to the gut and richly served by the same mesenteric lymphatic and venous drainage. But the microbiome hypothesis becomes more ambitious, and harder to defend, when it moves further from home. Multiple sclerosis is a disease of the central nervous system; the targets of the autoimmune attack are the myelin sheaths around nerve fibres in the brain and spinal cord. Rheumatoid arthritis attacks the synovial membranes of peripheral joints — wrists, hands, feet — tissues with no microbial community of their own. For the gut microbiome to participate in either disease requires a mechanism that reaches across the body.

The mechanism proposed, in both cases, is the same one we have been building toward: a gut dysbiosis that alters the production of short-chain fatty acids and other metabolites, that weakens the intestinal barrier and allows low-grade translocation of microbial products, and that together biases the systemic immune system toward inflammation and loss of peripheral tolerance. In the gut-adjacent organs, this bias expresses itself locally. In the brain and the joints, it expresses itself at a distance.

Multiple sclerosis: the Jangi cohort and its descendants

The first well-characterised gut microbiome study in multiple sclerosis was published in *Nature Communications* in 2016 by Sushrut Jangi, Howard Weiner, and colleagues at the Ann Romney Center for Neurologic Diseases at Brigham and Women's Hospital [REF:jangi2016]. The study was small by modern standards — 60 MS patients and 43 healthy controls, using 16S sequencing of stool — but it was carefully done, and it established several of the signal findings that subsequent work has confirmed.

The Jangi team found that patients with MS had significant alterations in gut microbial composition compared to controls: increased abundance of *Methanobrevibacter* (a methane-producing archaeon) and *Akkermansia muciniphila*, and decreased abundance of *Butyricimonas*, a butyrate-producing genus. The taxonomic differences correlated, in the same patients, with altered gene expression in circulating immune cells: genes involved in dendritic cell maturation, interferon signalling, and NF-κB signalling pathways in T cells and monocytes were differentially expressed between MS and control groups. A second independent cohort, studied for breath methane production, confirmed that elevated methanogen abundance translated into measurably elevated breath methane — the kind of whole-body functional confirmation that makes an association more believable.

The Jangi paper was a starting point rather than a conclusion. In the decade since, the MS microbiome field has grown into a substantial literature, but it has also hit the same methodological walls that troubled me in the T1D literature: small cohorts, heterogeneous methods, variable findings, and the replication difficulties that come with trying to capture a complex community with 16S sequencing. A 2023 systematic review by Ordoñez-Rodríguez and colleagues surveyed the field and concluded that while consistent themes had emerged — reduced butyrate producers, increased *Akkermansia*, altered Bacteroidetes/Firmicutes ratios — the between-study variation was substantial and the causal direction remained unsettled [REF:ordonezrodriguez2023].

The piece of work that most decisively advances the story is, in my view, a 2025 paper in *Cell Reports Medicine* by Luke Schwerdtfeger and colleagues, also from the Weiner laboratory at Brigham and Women's [REF:schwerdtfeger2025]. The Schwerdtfeger study did something the earlier literature had largely failed to do: it followed a prospective cohort of 114 MS subjects across three time points over two years, with paired microbiome and metabolome measurements, MRI brain imaging, and standardised disability and cognitive assessment. It is the first serious longitudinal dataset in which microbial changes can be related, patient by patient, to disease progression measured by objective imaging and clinical endpoints.

The central finding was this: patients whose gut communities showed higher abundance of SCFA-producing organisms — *Eubacterium hallii*, members of the Butyricoccaceae family, *Blautia* — had *slower* disease progression over the two-year follow-up, and the effect was concentrated in patients with progressive MS, the form of the disease in which conventional disease-modifying treatments

are least effective. The SCFA metabolomic signal paralleled the taxonomic one: higher faecal butyrate and propionate correlated with reduced MRI lesion burden and better disability scores.

This is, to my knowledge, the strongest longitudinal evidence to date that the gut microbiome is not just an epiphenomenon of MS but an active participant in its course. It does not prove causation. A trial in which SCFA-producing organisms, or SCFAs themselves, were administered to progressive MS patients and the course of their disease measured against placebo would be required for that. But it is the clearest signal yet that the variable clinical course of MS — why some patients progress and others do not — may be influenced by gut ecology, and it opens a plausible path to intervention. Butyrate-producing probiotic consortia and prebiotic dietary approaches are, as of 2026, in early clinical trials. Whether any will succeed remains to be seen.

The mechanism from gut to central nervous system

How does a gut community, sitting in the lumen of the colon, influence autoimmune attack on myelin sheaths in the spinal cord? Three mechanisms are in play, and all three are consistent with themes the reader has encountered in earlier chapters.

The first is *regulatory T cell induction*. Butyrate and propionate, produced by gut bacteria from dietary fibre, bind to the G-protein-coupled receptor GPR43 on immune cells and induce the differentiation of naive T cells into peripheral regulatory T cells — the T cells that actively suppress inappropriate immune responses, including responses to myelin. A gut that produces less butyrate produces fewer peripheral Tregs, and a body with fewer peripheral Tregs has a weaker brake on autoimmune responses wherever they occur. The effect is not MS-specific; it is a general bias of the immune system away from tolerance, and the disease that shows up in any given person is determined by their other genetic and environmental susceptibilities. In some susceptible people, it is MS. In others, it is one of the other diseases in this chapter.

The second is *barrier translocation*. The “leaky gut” mechanism described in the previous section applies here as well. Low-grade translocation of microbial products — LPS, bacterial DNA, muramyl dipeptide — produces chronic low-grade systemic inflammation, and a chronically inflamed peripheral immune system has a lower threshold for mounting attacks on previously tolerated self-antigens. The specifically neural targets of MS (myelin basic protein, myelin oligodendrocyte glycoprotein, proteolipid protein) are not themselves encountered by the gut immune system, but the systemic inflammatory tone in which T cells circulate affects how they respond to those antigens if they encounter them in the CNS.

The third — newest and least established — is *molecular mimicry* between microbial and myelin antigens. Several bacterial proteins share short sequence motifs with myelin components, and T cell clones isolated from MS lesions have occasionally been shown to cross-react with bacterial peptides. Whether this mechanism is a major contributor to MS pathogenesis or an interesting curiosity is still contested. It is worth mentioning because it raises the possibility that specific organisms, not just general community shifts, may matter for this disease.

Rheumatoid arthritis: the Scher paper and the *Prevotella* paradox

Rheumatoid arthritis has its own microbiome story, and it is, in some ways, the most instructive in this chapter — not because it is settled, but because the way it has failed to settle illustrates why microbiome research is harder than it looks.

The foundational paper is by Jose Scher and colleagues at New York University, published in *eLife* in 2013 [REF:scher2013]. The study sequenced the gut microbiomes of 114 people — a combination of new-onset untreated rheumatoid arthritis patients, chronic RA patients, patients with psoriatic arthritis, and healthy controls — and found a striking overrepresentation of the genus *Prevotella*, specifically the species *Prevotella copri*, in the new-onset untreated RA group. The effect was strong: *P. copri* accounted for a much larger share of the community in new-onset RA than in any other group, and its expansion was accompanied by a reduction in *Bacteroides* and several reportedly beneficial taxa.

To test whether the association was causal, the Scher team then did what microbiome researchers do when they want to move beyond correlation: they colonised mice. Germ-free mice given *P. copri* showed dominant colonisation and, when subsequently challenged with a chemically induced colitis model, more severe inflammation than controls. The interpretation offered in the paper was that *P. copri* might itself be a pro-inflammatory organism — a specific ecological driver of RA, rather than a marker of generic dysbiosis.

The Scher paper was hugely influential. *P. copri* rapidly entered the literature as a candidate RA-driving organism, and for a while it looked like RA might have found its AIEC. Then the replication work began, and the picture got complicated.

Some studies replicated the *P. copri* enrichment in RA. Others did not. Some found *P. copri* increased in pre-clinical RA, in healthy relatives of RA patients, and in other autoimmune conditions, suggesting the association was not RA-specific. And then, as the technology improved, it became clear that “*Prevotella copri*” was not really a single organism at all. It was a species label hiding a remarkable degree of strain diversity. The genome of one *P. copri* strain can differ from another by more than ten per cent of its gene content — an astonishing divergence for organisms that share a species name. A beautifully honest 2023 review in *Gut Microbes* by Abdelsalam and colleagues walked through the evidence and concluded, essentially, that the field had been measuring the wrong thing: species-level 16S profiling cannot resolve *P. copri* strains, and the clinical effects of one strain may bear little relation to the effects of another [REF:abdelsalam2023]. The reference-strain approach, dominant for a decade, had been misleading the field.

The fibre paradox

And then, in 2022, a paper by Jiang and colleagues in *Cell & Molecular Immunology* added the final and most surprising twist to the *Prevotella* story [REF:jiang2022]. The Jiang team colonised mice with *P. copri* and then fed them one of two diets: a standard laboratory diet, or a high-fibre diet of the kind that is routinely recommended for general gut health. They then induced arthritis using the collagen-induced arthritis (CIA) model, a standard mouse model of rheumatoid arthritis.

The result was the opposite of what most microbiome researchers would have predicted. In *P. copri*-colonised mice, the high-fibre diet did not protect against arthritis; it *worsened* it. Joint inflammation, histological damage, and arthritis severity scores were all significantly higher in *P. copri* + high-fibre mice than in *P. copri* + standard-diet mice, or than in germ-free controls on either diet. The mechanism, the authors showed, involved microbial metabolism: *P. copri* ferments fibre not primarily to butyrate and propionate but to *succinate* and *fumarate*, two metabolites that in appropriate amounts are normal intermediates of cellular metabolism but in excess promote pro-inflammatory macrophage polarisation. A gut community dominated by SCFA producers converts fibre into anti-inflammatory metabolites. A gut community dominated by *P. copri* converts the same fibre into pro-inflammatory metabolites. The substrate is the same. The output is different.

This is a lesson about context-dependence that deserves to be taken seriously, and it bridges cleanly to a theme we have been building throughout the book. In Volume 1 we argued that “good” and “bad” are misleading labels for microbes, because what matters is ecological context. Here is that principle in clinical form. Fibre is not unambiguously good. It is good for hosts whose gut communities are equipped to ferment it into anti-inflammatory products. It can be harmful for hosts whose gut communities are equipped to ferment it into pro-inflammatory products. The general advice to “eat more fibre,” which I gave with some confidence in Chapter 20 and which remains broadly correct for the population, can fail badly for individuals whose ecology produces an unusual metabolic profile. In a genuinely precision-medicine future, “should you eat more fibre?” would be a question to answer with reference to the patient’s gut microbiome, not to a population-level recommendation.

What does the *P. copri*/fibre story mean for RA patients today? The honest answer is that it does not change clinical practice, at least not yet. RA remains a disease managed with disease-modifying antirheumatic drugs (methotrexate, sulfasalazine) and biologics (TNF inhibitors, JAK inhibitors), and microbiome-directed therapies for RA remain experimental. But the story does mean that the one-size-fits-all dietary advice historically given to RA patients — the “anti-inflammatory diet” heavy in whole grains and vegetables — may work well for some patients and poorly for others, and the difference may lie in which *Prevotella* strains they carry. Until we have clinical tests that can resolve strain-level variation cheaply and in real time, the best we can do is to be open to individual variation in response and to avoid dogmatic advice.

The broader lesson for the chapter is this: the microbiome influences autoimmune disease in ways that are more context-dependent than the early literature suggested. The same species can be protective or harmful depending on strain. The same dietary intervention can reduce or exacerbate inflammation depending on the community metabolising it. The same mechanistic principle — dysbiosis $\square$ altered metabolism $\square$ altered immune tone — works in both directions, and the direction in any individual patient is determined by who is in their gut and what those residents are doing. This is harder to translate into clinical practice than a simpler story would be, but it is closer to the truth of what is going on.

26.5 Coeliac Disease: Genes, Gluten, and Microbes

Of the five diseases in this chapter, coeliac disease has the cleanest experimental set-up. The trigger is known, isolable, and removable. Gluten — specifically the gliadin fraction of the wheat protein, and its homologues in rye and barley — is the environmental factor that, in genetically susceptible individuals, drives the autoimmune attack on the small intestinal mucosa. Remove gluten from the diet, and the disease remits (mostly). Reintroduce gluten, and it returns. No other autoimmune disease in the clinic has such a clean trigger-and-remove structure, and coeliac disease is therefore an unusually good laboratory for asking the question this chapter keeps returning to: what does the microbiome add?

The answer, it turns out, is more than you might expect. And the story is a useful corrective to the idea that once you know the trigger and the HLA type, you have explained the disease.

The HLA paradox

Coeliac disease requires a specific genetic background: the HLA-DQ2 or HLA-DQ8 haplotypes, which encode particular MHC class II molecules able to present deamidated gliadin peptides to

CD4+ T cells. This is a very strong genetic requirement — essentially every coeliac patient carries DQ2 or DQ8 — and it has been the foundation of coeliac immunology for decades.

The paradox is that this genetic requirement is extraordinarily *common*. Somewhere between thirty and forty per cent of the general population in European-descended communities carries DQ2 or DQ8. Coeliac disease, by contrast, affects roughly one per cent of the same populations. Something, in other words, separates the ninety-nine per cent of HLA-DQ2/8 carriers who eat gluten their whole lives without incident from the one per cent who develop the disease. Gluten exposure is common. The HLA type is common. The co-occurrence of both is common. But manifest coeliac disease is not. Whatever determines who tips over from potential to actual disease is important, and it is not the HLA genotype (because that is shared with the ninety-nine per cent) and it is not the exposure to gluten (ditto).

The microbiome hypothesis proposes that the shift from “potential coeliac” to “manifest coeliac” is a second hit delivered by the gut community. Children with the HLA-DQ2/8 background who also have a particular early-life microbiome configuration may be pushed over the diagnostic threshold, while those with the same HLA background and a different microbiome may pass their whole lives in a subclinical state of elevated risk without ever tipping into disease.

This is, of course, a testable prediction, and it has been tested.

The dysbiotic signature of coeliac disease

A 2024 review by Giulia Catassi and colleagues in *Best Practice & Research Clinical Gastroenterology* is the most recent comprehensive synthesis of the coeliac microbiome literature [REF:catassi2024]. The review pools the findings of dozens of case-control and longitudinal studies, and the consistent pattern it identifies across cohorts is this: coeliac patients have reduced populations of *Bifidobacterium* and *Lactobacillus*, and expanded populations of *Bacteroides*, *Escherichia coli*, and other Gram-negative organisms.

This is a shift with mechanistic teeth. *Bifidobacterium* and *Lactobacillus* species are among the main producers of short-chain fatty acids and of specific lactic acid derivatives that support epithelial barrier function and regulatory T cell tone. Their depletion is a depletion of exactly the ecological services that a genetically susceptible mucosa would most benefit from. The expansion of Gram-negative organisms brings with it increased LPS exposure and an innate immune milieu that is, on balance, more inflammatory. A gut community shifted in this direction is one in which gliadin exposure — chemically modified by the tissue enzyme transglutaminase-2 to produce the deamidated peptides that HLA-DQ2 and DQ8 preferentially present — encounters a more permissive and more inflammatory local environment than it would in a gut with robust *Bifidobacterium* populations.

The interaction with early-life factors is, if anything, more interesting than the steady-state signature. Caesarean delivery, formula feeding, and early antibiotic exposure — all of which shift the infant gut community away from the *Bifidobacterium*-dominated state we saw in the TEDDY data — have each been associated with increased coeliac disease risk in HLA-susceptible children. The biggest risk factor in some studies is the combination: a DQ2/8-positive child born by Caesarean, formula-fed, and exposed to antibiotics in the first year of life has a substantially higher coeliac risk than a DQ2/8-positive child with none of those exposures. The pattern is identical to the one we saw in T1D, in IBD, and in several other conditions in this chapter, and it reinforces the common-mechanism idea.

The gluten-free diet does not fully normalise the microbiome

One of the most surprising findings in the coeliac microbiome literature is that a strict gluten-free diet, while it dramatically improves clinical symptoms and heals the villous atrophy that defines the disease, does *not* fully restore the gut microbiome to a non-coeliac configuration. Even after years on a strict GFD, coeliac patients retain a distinct dysbiotic signature: lower diversity, lower *Bifidobacterium*, and altered *Bacteroides*/Firmicutes ratios compared to matched controls.

Three explanations are on the table, and they are not mutually exclusive. The first is that the GFD itself perturbs the microbiome. A gluten-free diet, in practice, tends to be lower in whole-grain fibre and in certain resistant starches than the diet it replaces, because many of the cheapest and most convenient gluten-free products are made from refined rice flour, tapioca, and potato starch. The dysbiosis may, in part, be a diet-induced dysbiosis overlaid on a disease-induced one. The second is that the gut epithelium in long-term coeliac patients remains subtly abnormal even when the villi look normal on biopsy — transcriptomic studies continue to show altered epithelial gene expression in treated coeliac disease — and the persistent mucosal differences sustain a persistent community shift. The third is that the infant microbiome was established before the disease was ever recognised, and the original ecological trajectory, once set, is hard to reverse. The founding community shapes the adult community, and fixing the downstream disease does not retroactively fix the upstream ecology.

A 2022 Saudi paediatric cohort by El Mouzan and colleagues adds a useful non-Western data point [REF:elmouzan2022]. The study enrolled forty-five children with newly diagnosed coeliac disease in Saudi Arabia, and compared their microbiomes to those of healthy local controls. Several of the “classic” coeliac findings — reduced *Bifidobacterium*, altered *Bacteroides* — replicated, but with regional twists, reminding us that microbiome signatures are population-dependent and that patterns identified in one cohort do not automatically generalise to another. This is a point I keep making in this chapter, and it keeps being relevant: the microbiome is local, and inferring causation from patterns observed in one population and then generalising to another is a methodological move that requires caution.

Probiotics for coeliac disease: the honest answer

Commercial interest in probiotic treatments for coeliac disease is substantial. The logic is superficially appealing — if *Bifidobacterium* and *Lactobacillus* are depleted, restoring them might help — and the product category is already well established, so the transition to a new indication is commercially cheap.

The evidence, as of 2026, is modest. Randomised trials of specific probiotic strains in coeliac disease have shown two kinds of results. A handful have shown reductions in gastrointestinal symptoms (bloating, diarrhoea) during gluten challenge, suggesting that probiotics may help with symptom management. Almost none have shown effects on mucosal healing, on serological markers of disease activity, or on the fundamental autoimmune process. The picture is consistent with probiotics doing what they often do elsewhere in this book: delivering modest symptomatic benefit for some patients without fundamentally altering the underlying disease.

The most honest current statement is that probiotics are not a treatment for coeliac disease. A strict gluten-free diet remains the only evidence-based therapy. Probiotics may, in specific strains and specific contexts, help with residual symptoms in some patients. They do not replace the diet, they do not accelerate mucosal healing in any reproducible way, and they certainly do not allow gluten

reintroduction. The patient who is told otherwise — and this phrasing creeps into the marketing of many products — is being misled.

Towards an integrated view

Coeliac disease is an especially clear example of the general model this chapter has been developing. There is a genetic predisposition (HLA-DQ2/8) that is common, and a dietary trigger (gluten) that is also common. The microbiome appears to determine, in part, who among the susceptible tips over into disease. The mechanism involves altered SCFA production, barrier permeability, and the local inflammatory environment in which gluten peptides are presented to T cells. Early-life exposures (delivery mode, feeding, antibiotics) shape the microbiome in ways that correlate with later coeliac risk. The disease, once established, produces further ecological shifts that are only partly corrected by the standard treatment. And the therapeutic opportunities beyond the gluten-free diet remain largely unproven, despite substantial commercial enthusiasm.

Put this alongside the IBD, T1D, MS, and RA stories and the outlines of a common picture emerge: a disturbed gut community, a weakened barrier, altered microbial metabolism, and a bias of the immune system toward inflammation and loss of tolerance, expressed in different tissues according to the patient's other susceptibilities. The outlines are clearer than they were a decade ago. What remains to be filled in is the line between the outlines and clinical practice — the intervention trials, the biomarkers, the personalisation tools — that would turn a compelling model into a treatment a patient can hold in their hand.

26.6 A Common Mechanism or a Convergent Phenotype?

Five diseases, five stories, one recurrent set of themes: reduced microbial diversity, depleted short-chain fatty acid producers, increased intestinal permeability, low-grade translocation of microbial products, altered regulation of T cell tolerance, and clinical disease expressed in tissues that reflect the patient's genetic susceptibilities. It is tempting, after running through the five case studies, to declare that we have identified *the* mechanism of autoimmune disease — a single pathological cascade with different tissue targets — and to begin writing prescriptions.

That declaration would be premature. The data are consistent with a common mechanism. They are also consistent with several different mechanisms that happen to converge on similar-looking endpoints. Deciding between these interpretations matters, because it determines whether microbiome-directed interventions should be expected to help across the whole family of autoimmune diseases or whether each disease will require its own specific ecological fix. Let me lay out the two positions as fairly as I can.

The “common pathway” view

In its strongest form, the common-pathway view holds that autoimmune disease is, substantially, a disease of lost microbial regulation. The twentieth-century rise in autoimmunity tracks the twentieth-century disruption of the human microbiome, and the mechanistic core is the same in every case: loss of the SCFA-producing, Treg-inducing, barrier-supporting community that our immune system evolved to rely on. The genetic predisposition determines *which* tissue is attacked when tolerance fails; the microbial loss determines *whether* tolerance fails at all.

This view is supported by the observation that autoimmune diseases cluster in individuals and in families. A person with one autoimmune diagnosis is substantially more likely than the general population to develop a second. Coeliac disease raises the risk of T1D; T1D raises the risk of autoimmune thyroid disease; IBD raises the risk of primary sclerosing cholangitis, ankylosing spondylitis, and rheumatoid arthritis. If each disease were caused by a wholly distinct mechanism, this clustering would be puzzling. If they share a common immunological substrate — a generally inflammatory, tolerance-impaired immune system — the clustering is expected.

It is also supported by the remarkable consistency of the SCFA theme. Across every disease in this chapter, and in the broader literature, the finding that SCFA producers are depleted and their products reduced is one of the most reproducible in the field. Kim's 2023 integrative review in *Cell & Molecular Immunology* covers T1D, MS, RA, and systemic lupus erythematosus and shows that all four are associated with reduced butyrate-producing capacity, altered SCFA serum profiles, and impaired Treg induction by gut-derived metabolites [REF:kim2023]. If the SCFA story is the core of the model, and if SCFA production is a general microbial function rather than a disease-specific one, then autoimmune diseases are indeed variations on a common theme.

The “convergent phenotype” view

The alternative view accepts the observations but disputes the interpretation. On this reading, the common patterns across autoimmune diseases are not evidence of a shared mechanism; they are evidence that inflammation has a limited repertoire. Whatever perturbs the gut — whether it is Crohn's-associated AIEC, T1D-associated early antibiotic exposure, MS-associated loss of butyrate producers, RA-associated *P. copri* expansion, or coeliac-associated *Bifidobacterium* depletion — produces, via different upstream routes, a similar-looking downstream state: low diversity, low SCFAs, leaky barrier, inflammatory tone. The shared endpoint reflects the limited number of ways a gut community can malfunction, not a shared cause of the malfunction.

This view is harder to test than the first, because the observable patterns are largely the same under both interpretations. But it is supported by several things. The genetic architectures of the five diseases are substantially different from each other: IBD is driven heavily by variants in innate immune and autophagy genes; T1D and coeliac disease by HLA and other adaptive-immune loci; MS by a distinct set of MHC and cytokine pathway variants; RA by a combination of HLA and post-translational modification machinery. If the diseases shared a single microbial mechanism, we would expect more overlap in susceptibility loci than we in fact observe. The specific microbial signals also differ substantially between diseases — *Faecalibacterium* loss in IBD, *P. copri* in RA, *Methanobrevibacter* in MS, *Candida* in IBD's mycobiome — and the specifics are not fully compatible with a single causal mechanism. A common downstream state, reached by several upstream routes, fits the data at least as well as a shared mechanism.

What this means for intervention

The practical distinction between the two views matters because it predicts how microbiome interventions should be designed and how broadly they should apply.

If the common-pathway view is correct, a generic intervention that restores SCFA-producing capacity — butyrate producer consortia, dietary fibre optimised for the individual's microbiome, next-generation probiotics selected for butyrogenic properties — might be expected to help across the whole range of autoimmune diseases. The approach would be unified, and the per-disease tailoring would be modest.

If the convergent-phenotype view is correct, a common downstream intervention might help symptomatically without addressing the specific upstream drivers, and each disease might require its own distinct intervention strategy. Crohn's disease might need phages targeting AIEC plus antifungals. RA might need strain-specific *P. copri* replacement. T1D might need early-life interventions directed at the infant microbiome specifically. MS might need butyrate consortia tuned for CNS effects. The approach would be disease-specific, the per-disease tailoring would be substantial, and generic “anti-autoimmune probiotics” would be unlikely to succeed.

As of 2026, the evidence base does not firmly distinguish between these positions, and in my view the answer is likely to be “both, to different degrees in different diseases.” The common downstream state is real and almost certainly contributes to the clinical phenotype. The specific upstream drivers are also real, and in some diseases (Crohn's, possibly RA) they are likely to be more important than the shared downstream features. The successful therapies of the next decade will probably need to address both levels: restoring the general ecological functions that have been lost, *and* correcting the specific disease-associated shifts that drive the particular tissue target. A 2022 review by Belvončíková, Maronek and Gardlík surveyed the early and still-preliminary attempts to use faecal microbiota transplantation to correct dysbiosis across this whole range of autoimmune conditions, and drew much the same conclusion: FMT produces signal in some diseases (notably ulcerative colitis, as Chapter 24 discussed), tentative signal in others, and near-silence in several, in a pattern that is more compatible with multiple disease-specific mechanisms than with a single unifying one [REF:belvoncikova2022].

Caution, not despair

I want to end this section, and this chapter, with a reflection on what a realistic and honest microbiome medicine for autoimmune disease would look like, because I think the field is due for a recalibration.

The hype cycle has been, in places, frankly absurd. Every few months, a popular-press piece announces that scientists have “found the bacterium that causes” multiple sclerosis, or Crohn's disease, or type 1 diabetes. The piece is almost always based on a single paper with a small cohort and a modest effect size. The conclusion is almost always overstated. The commercial follow-on — a probiotic marketed as targeting that very bacterium, or a fermented product claimed to correct the dysbiosis — almost always lacks any evidence of efficacy in the disease in question. The cycle has done real harm to real patients, who spend real money on real products that do real nothing, and who then lose faith in microbiome science as a whole when the products fail.

Against that backdrop, the genuine science this chapter has described is both more modest and more interesting than the hype version. The microbiome is part of the story of autoimmune disease, not the whole story. The mechanisms are partially understood, partially plausible, and partially still speculative. The interventions that work — for specific indications, in specific patients, at specific doses — are few, and most of the interesting ones are not yet approved as treatments. The interventions that are *marketed* are, with a few exceptions, not the ones that have the evidence behind them.

A patient with Crohn's disease in 2026 is still best managed with biologics, immunomodulators, and sometimes surgery. A patient with T1D still requires insulin. A patient with MS benefits from disease-modifying therapies, not from probiotic capsules. A patient with RA takes methotrexate and, when needed, a TNF inhibitor. A patient with coeliac disease eats gluten-free. The microbiome may, in the next decade or two, add adjunctive tools to these primary treatments — SCFA-producing

consortia, personalised dietary strategies, targeted cross-kingdom interventions — but it will not replace them, and it should not be pretended that it already has.

What the microbiome model does offer, right now, is explanatory power. It tells us why these diseases cluster, why they are rising in parallel, why early-life exposures matter, why some children develop disease and their identical twins do not, and why genetics alone cannot predict who will get sick. An explanation is not a treatment. But it is a foundation for the treatments that will come, and the foundation is, I think, largely sound.

The next chapter turns to another group of diseases in which this same ecological framework has recently gained substantial traction: the metabolic diseases. Obesity, type 2 diabetes, non-alcoholic fatty liver disease — conditions that were, until very recently, thought of as problems of caloric balance and endocrine signalling — turn out, on closer inspection, to have a microbial dimension too. The leaky-gut axis we have just traced through the autoimmune diseases appears, in rearranged form, in the metabolic ones, and the cross-kingdom thinking that resolved ASCA after thirty years is now being brought to bear on the gut-liver axis and on the role of the microbiome in energy harvest. The story there is more advanced in some ways than it is here — metformin's microbial mechanism, for instance, is better established than anything in this chapter — but it is not fundamentally different. The body is an ecosystem. The diseases of modern life are, in many cases, ecosystem malfunctions. The rest of this book is about making that case, disease by disease.

Key Chapter References

[REF:iliev2025] Iliev ID, Ananthakrishnan AN & Guo CJ (2025). Microbiota in inflammatory bowel disease: mechanisms of disease and therapeutic opportunities. *Nat Rev Microbiol.* 23(8):509–524. doi:10.1038/s41579-025-01163-0

[REF:gevers2014] Gevers D, Kugathasan S, Denson LA, et al. (2014). The treatment-naïve microbiome in new-onset Crohn's disease. *Cell Host Microbe.* 15(3):382–392. doi:10.1016/j.chom.2014.02.005

[REF:haneishi2023] Haneishi Y, Furuya Y, Hasegawa M, et al. (2023). Inflammatory Bowel Diseases and Gut Microbiota. *Int J Mol Sci.* 24(4):3817. doi:10.3390/ijms24043817

[REF:darfeuille2004] Darfeuille-Michaud A, Boudeau J, Bulois P, et al. (2004). High prevalence of adherent-invasive *Escherichia coli* associated with ileal mucosa in Crohn's disease. *Gastroenterology.* 127(2):412–421. doi:10.1053/j.gastro.2004.04.061

[REF:henke2019] Henke MT, Kenny DJ, Cassilly CD, Vlamakis H, Xavier RJ, Clardy J (2019). *Ruminococcus gnavus*, a member of the human gut microbiome associated with Crohn's disease, produces an inflammatory polysaccharide. *Proc Natl Acad Sci USA.* 116(26):12672–12677. doi:10.1073/pnas.1904099116

[REF:hugot2001] Hugot JP, Chamaillard M, Zouali H, et al. (2001). Association of NOD2 leucine-rich repeat variants with susceptibility to Crohn's disease. *Nature.* 411(6837):599–603. doi:10.1038/35079107

[REF:leibovitzh2024] Leibovitzh H, Nayeri S, Borowski K, et al. (2024). Inflammatory bowel disease associated with primary sclerosing cholangitis is associated with an altered gut microbiome and bile acid profile. *J Crohns Colitis.* 18(12):1957–1966. doi:10.1093/ecco-jcc/jjae096

- [REF:sendid2023] Sendid B, Cornu M, Cordier C, Bouckaert J, Colombel JF, Poulain D (2023). From ASCA breakthrough in Crohn's disease and *Candida albicans* research to thirty years of investigations about their meaning in human health. *Autoimmun Rev.* 23(2):103486. doi:10.1016/j.autrev.2023.103486
- [REF:fuchtbauer2025] Füchtbauer JD, Berentsen EM, Lindhardt CL, et al. (2025). Gut Mycobiome in Inflammatory Bowel Disease: A Systematic Review of Case-Control Studies. *Inflamm Bowel Dis.* 31(11):3213–3224. doi:10.1093/ibd/izaf217
- [REF:hsu2023] Hsu C, Ghannoum M, Cominelli F, Di Martino L (2023). Mycobiome and Inflammatory Bowel Disease: Role in Disease Pathogenesis, Current Approaches and Novel Nutritional-based Therapies. *Inflamm Bowel Dis.* 29(3):470–479. doi:10.1093/ibd/izac156
- [REF:shao2022] Shao TY, Haslam DB, Bennett RJ, Way SS (2022). Friendly fungi: symbiosis with commensal *Candida albicans*. *Trends Immunol.* 43(9):706–717. doi:10.1016/j.it.2022.07.003
- [REF:fasano2012] Fasano A (2012). Leaky gut and autoimmune diseases. *Clin Rev Allergy Immunol.* 42(1):71–78. doi:10.1007/s12016-011-8291-x
- [REF:divincenzo2023] Di Vincenzo F, Del Gaudio A, Petito V, Lopetuso LR, Scaldaferri F (2023). Gut microbiota, intestinal permeability, and systemic inflammation: a narrative review. *Intern Emerg Med.* 19(2):275–293. doi:10.1007/s11739-023-03374-w
- [REF:mu2017] Mu Q, Kirby J, Reilly CM, Luo XM (2017). Leaky Gut As a Danger Signal for Autoimmune Diseases. *Front Immunol.* 8:598. doi:10.3389/fimmu.2017.00598
- [REF:vatanen2018] Vatanen T, Franzosa EA, Schwager R, et al. (2018). The human gut microbiome in early-onset type 1 diabetes from the TEDDY study. *Nature.* 562(7728):589–594. doi:10.1038/s41586-018-0620-2
- [REF:stewart2018teddy] Stewart CJ, Ajami NJ, O'Brien JL, et al. (2018). Temporal development of the gut microbiome in early childhood from the TEDDY study. *Nature.* 562(7728):583–588. doi:10.1038/s41586-018-0617-x
- [REF:luppi2024] Luppi S, Aldegheri L, Azzalini E, et al. (2024). Unravelling the Role of Gut and Oral Microbiota in the Pediatric Population with Type 1 Diabetes Mellitus. *Int J Mol Sci.* 25(19):10611. doi:10.3390/ijms251910611
- [REF:asantebaadu2025] Asante Baadu F, Yalley C, Ohene Yeboah AR, et al. (2025). Microbiome Imbalance and Pediatric Type 1 Diabetes Mellitus: An Updated Systematic Review of Gut Dysbiosis Evidence. *Cureus.* 17(8):e89279. doi:10.7759/cureus.89279
- [REF:jangi2016] Jangi S, Gandhi R, Cox LM, et al. (2016). Alterations of the human gut microbiome in multiple sclerosis. *Nat Commun.* 7:12015. doi:10.1038/ncomms12015
- [REF:ordonezrodriguez2023] Ordoñez-Rodríguez A, Roman P, Rueda-Ruzafa L, Campos-Rios A, Cardona D (2023). Changes in Gut Microbiota and Multiple Sclerosis: A Systematic Review. *Int J Environ Res Public Health.* 20(5):4624. doi:10.3390/ijerph20054624
- [REF:schwerdtfeger2025] Schwerdtfeger LA, Jangi S, Yeh WZ, et al. (2025). Gut microbiota and metabolites are linked to disease progression in multiple sclerosis. *Cell Rep Med.* 6(4):102055. doi:10.1016/j.xcrm.2025.102055
- [REF:scher2013] Scher JU, Sczesnak A, Longman RS, et al. (2013). Expansion of intestinal *Prevotella copri* correlates with enhanced susceptibility to arthritis. *eLife.* 2:e01202.

doi:10.7554/eLife.01202

[REF:jiang2022] Jiang L, Shang M, Yu S, et al. (2022). A high-fiber diet synergizes with *Prevotella copri* and exacerbates rheumatoid arthritis. *Cell Mol Immunol.* 19(12):1414–1424. doi:10.1038/s41423-022-00934-6

[REF:abdelsalam2023] Abdelsalam NA, Hegazy SM, Aziz RK (2023). The curious case of *Prevotella copri*. *Gut Microbes.* 15(2):2249152. doi:10.1080/19490976.2023.2249152

[REF:kim2023] Kim CH (2023). Complex regulatory effects of gut microbial short-chain fatty acids on immune tolerance and autoimmunity. *Cell Mol Immunol.* 20(4):341–350. doi:10.1038/s41423-023-00987-1

[REF:catassi2024] Catassi G, Aloï M, Giorgio V, Gasbarrini A, Cammarota G, Ianiro G (2024). The role of microbiome in the development of gluten-related disorders. *Best Pract Res Clin Gastroenterol.* 72:101951. doi:10.1016/j.bpg.2024.101951

[REF:elmouzan2022] El Mouzan M, Assiri A, Al Sarkhy A, et al. (2022). Microbiota profile of new-onset celiac disease in children in Saudi Arabia. *Gut Pathog.* 14(1):37. doi:10.1186/s13099-022-00493-1

[REF:belvonicikova2022] Belvonicikova P, Maronek M, Gardlik R (2022). Gut Dysbiosis and Fecal Microbiota Transplantation in Autoimmune Diseases. *Int J Mol Sci.* 23(18):10729. doi:10.3390/ijms231810729

[REF:aleman2023] Aleman RS, Moncada M, Aryana KJ (2023). Leaky Gut and the Ingredients That Help Treat It: A Review. *Molecules.* 28(2):619. doi:10.3390/molecules28020619

Chapter 27: Metabolic Disease and Obesity

27.0 The Weight of the Ecosystem

In 2024, a consortium of researchers published data in *The Lancet* confirming what public health officials had long suspected: more than one billion people on the planet were now living with obesity. The number had roughly quadrupled since 1990. In some countries — the United States, the Gulf States, several Pacific island nations — more than forty per cent of the adult population met the clinical threshold. And the trajectory, despite decades of public health campaigns, dietary guidelines, and pharmaceutical interventions, showed no sign of bending.

The conventional explanation for obesity is thermodynamic. You eat more energy than you expend, and the surplus is stored as fat. This is true as far as it goes — the first law of thermodynamics is not negotiable — but it has always been an incomplete account, because it leaves the most important clinical question unanswered: *why do different people, eating very similar diets and living very similar lives, gain very different amounts of weight?*

The standard answers to that question have been genetics and behaviour. Both are real. Genome-wide association studies have identified hundreds of loci associated with body mass index, and the heritability of BMI is estimated at sixty to seventy per cent — among the highest of any common trait. But the obesity epidemic has unfolded in two generations, and the genome has not changed. Something in the environment is interacting with those genetic susceptibilities to push a growing fraction of the population past the threshold into clinical disease, and the identity of that environmental factor — or, more honestly, that constellation of environmental factors — is the central question of modern obesity research.

Diet has changed. Physical activity has declined. Sleep is shorter. Stress is chronic. Ultra-processed foods have colonised the food supply. All of these contribute. But over the past two decades, a further variable has entered the equation, one that this book has been building toward from its first page: the gut microbiome.

The proposition is not that the microbiome *causes* obesity. That claim would be too strong and the evidence does not support it. The proposition, rather, is that the microbiome is a mediator — a biological system that sits between what you eat and what your body does with it, and whose composition and functional state influence how much energy you extract from food, how you store fat, how your liver handles lipids, how your pancreas manages blood sugar, and how your brain registers satiety. Individual differences in that mediating system help explain why two people can eat the same meal and emerge with different metabolic consequences. And systematic shifts in that system — driven by the same dietary, pharmaceutical, and environmental changes that have reshaped modern life — may be one of the reasons the metabolic disease curve keeps rising.

This chapter examines that proposition across three metabolic conditions: obesity itself, type 2 diabetes, and the cluster of cardiovascular and hepatic diseases that travel with them. The evidence is, in places, strong. In others, it is suggestive but incomplete. The microbiome's role in metabolic disease is better established, mechanistically, than its role in most of the autoimmune conditions

we examined in the previous chapter — the pathways from dysbiosis to insulin resistance, from microbial metabolites to hepatic fat accumulation, from gut-derived molecules to arterial plaque are traceable in molecular detail. But the gap between mechanistic understanding and clinical application is, here as elsewhere, wider than the popular press tends to suggest. No microbiome-based therapy has yet produced clinically meaningful weight loss in humans. The most effective anti-obesity drugs of the current era — the GLP-1 receptor agonists — were not developed from microbiome science, even though, as we will see, they turn out to interact with the microbiome in ways that are only now being understood.

The chapter that follows tries to hold both of these truths in view. The microbiome matters for metabolic disease. It does not yet offer a treatment for it. Understanding why is, I think, as instructive as the science itself.

27.1 The Famous “Obese Microbiome” Experiments — What They Showed and What They Didn’t

The modern field of obesity-microbiome research began in a laboratory at Washington University in St Louis, in the mid-2000s, with a series of experiments that are now taught in every microbiology graduate programme on earth. They were conducted by Jeffrey Gordon’s group — the same laboratory that had pioneered the use of germ-free mice to study host-microbe interactions — and they asked a question so simple that it is surprising nobody had asked it before: does the gut microbiome of an obese animal look different from that of a lean one?

The founding observation

In 2005, Ruth Ley, Fredrik Bäckhed, Peter Turnbaugh, and Gordon published the first systematic comparison of gut microbial communities in obese and lean mice [REF:ley2005]. They used genetically obese *ob/ob* mice — animals carrying a mutation in the *leptin* gene that makes them chronically hungry and profoundly fat — and compared their cecal microbiota with that of lean *ob/+* and wild-type siblings, all fed the same polysaccharide-rich diet. The comparison was done using 16S ribosomal RNA gene sequencing, the census tool we met in Chapter 4.

The result was striking. At the phylum level — the broadest meaningful division in bacterial taxonomy — the obese mice had a roughly fifty per cent reduction in *Bacteroidetes* and a proportional increase in *Firmicutes* compared with their lean siblings. The shift was consistent, statistically robust, and independent of kinship. Something about the obese state, whether caused by the obesity itself or contributing to it, was associated with a wholesale restructuring of the gut microbial community at its most fundamental taxonomic level.

A year later, Ley and colleagues extended the observation to humans [REF:ley2006]. Twelve obese volunteers were placed on either a fat-restricted or a carbohydrate-restricted low-calorie diet and followed for a year. At baseline, the obese subjects had a lower proportion of *Bacteroidetes* and a higher proportion of *Firmicutes* than lean controls. Over the course of weight loss, the *Bacteroidetes* proportion increased, regardless of which diet was used — the shift tracked with weight loss itself, not with the macronutrient composition of the diet.

These two papers established the *Firmicutes*-to-*Bacteroidetes* ratio as the headline metric of the obese microbiome. For several years, the ratio was cited in virtually every review, every conference keynote, and every popular-press article about the gut and weight. It was a clean, memorable result

— two phyla, one going up, one going down, tracking with body weight — and it had the advantage of being easy to explain.

The transmissibility experiment

But the really consequential paper came in December 2006, when Turnbaugh, Ley, and Gordon published a study in *Nature* that moved the field from association to mechanism [REF:turnbaugh2006]. The experiment was elegant. They took gut microbiota from genetically obese *ob/ob* mice and from lean controls and transplanted them into germ-free recipient mice — animals that had been raised without any microbes of their own. Both groups of recipients were fed the same diet. And the recipients that received the “obese” microbiota gained significantly more body fat than those that received the “lean” microbiota.

The difference was not enormous — about two percentage points of body fat over two weeks — but it was consistent, and it demonstrated something no observational study could: that the microbial community associated with obesity was not merely a passive consequence of the obese state. It was, at least in part, *functional*. It could transmit a metabolic phenotype to a new host.

The mechanism, revealed by metagenomic analysis of the transplanted communities, was increased energy harvest. The obese microbiota contained more genes involved in the breakdown of otherwise indigestible dietary polysaccharides — complex plant fibres that the host’s own enzymes cannot process. These genes enabled the community to extract more calories from the same food, producing short-chain fatty acids that the host could absorb and use. The obese microbiota was, in effect, a more efficient engine for converting dietary fibre into usable energy.

Think of it this way. Two cars — same make, same model — are filled with the same fuel. But one has been fitted with a turbocharger that extracts a few extra per cent of energy from each litre. Over time, that small difference in efficiency adds up. The “turbocharger,” in this analogy, is the microbial community, and its efficiency is determined not by a single species but by the collective metabolic capacity of the entire ecosystem.

The correction

These experiments were landmark science. They opened a field, launched a thousand studies, and changed the way both researchers and the public thought about obesity. But they also, inevitably, generated an oversimplified narrative — one that the subsequent two decades of research have substantially complicated.

The *Firmicutes*-to-*Bacteroidetes* ratio, the headline finding of the original papers, has not replicated consistently in humans. A 2014 meta-analysis, and several large cohort studies since, have found that the ratio is not a reliable discriminator between obese and lean individuals across different populations, different diets, and different sequencing methodologies. Some studies find it; others do not. The reasons are multiple: the two phyla are enormous and internally diverse (*Firmicutes* alone contains thousands of species with wildly different metabolic capabilities), the ratio is sensitive to methodological choices (DNA extraction, primer selection, bioinformatics pipeline), and human obesity is far more heterogeneous than the *ob/ob* mouse model.

The field has accordingly moved beyond phylum-level ratios toward more granular and more functional measures. The most influential reframing came from the MetaHIT consortium’s 2013 study of 292 Danish adults, led by Emmanuelle Le Chatelier and colleagues [REF:lechatelier2013].

Rather than asking which phyla were present, they asked how many microbial genes each person carried — a measure of the community's total functional capacity, which they called “gene richness.” They found that the study population divided into two groups: a “high gene count” group with diverse, functionally rich communities, and a “low gene count” group with depleted, functionally impoverished ones. The low-gene-count individuals — about twenty-three per cent of the population — had more adiposity, more insulin resistance, more dyslipidaemia, and more systemic inflammation than the high-gene-count group. And among obese individuals, those with low gene counts gained more weight over time.

This was a more nuanced and more powerful finding than the *Firmicutes/Bacteroidetes* ratio. It said that what matters for metabolic health is not which broad group of bacteria you carry, but how rich and functionally diverse your community is — and that a depleted community, regardless of its taxonomic composition, is a metabolically vulnerable one. It echoed the diversity-health association we have seen throughout this book, from the developing infant microbiome (Chapter 3) to the IBD microbiome (Chapter 26), and it set the stage for a more ecological understanding of the obesity-microbiome relationship.

A 2026 review by Davide Masi, Mikiko Watanabe, and Karine Clément in *Cell Reports Medicine* provides the most current synthesis of this evolution, and it is the paper I would recommend to any reader wanting a single authoritative entry point into the field as it stands today [REF:masi2026]. Its central argument is that the obesity microbiome is best understood not through simple taxonomic signatures but through “metabolic phenotype stratification” — identifying microbiome-driven metabolic profiles that predict individual responses to dietary, surgical, and pharmacological interventions. The era of “obese microbiome versus lean microbiome” is over. The era of precision metabolic microbiology is, tentatively, beginning.

27.2 Energy Harvest, Bile Acids, and Appetite Signalling

The Turnbaugh experiment established the principle: an obese microbiota can extract more energy from the same diet. But a principle is not a mechanism. Over the two decades since that paper, researchers have mapped at least three distinct pathways through which the gut microbiome influences metabolic energy balance. Each of them intersects with biology we have already met in this book, and each of them has clinical implications that extend well beyond body weight.

The SCFA pathway

The first and most straightforward pathway is the one Turnbaugh identified: enhanced fermentation of dietary fibre into short-chain fatty acids (SCFAs) — acetate, propionate, and butyrate — which the host absorbs and metabolises. We covered the production and functions of SCFAs in Chapter 20, and their role in immune regulation in Chapter 26. Here, the relevant question is energetic: how much do SCFAs contribute to the host's total energy intake?

The answer is modest but not negligible. In a typical Western diet, colonic fermentation of fibre produces an estimated 200 to 600 kilocalories per day in SCFAs, depending on fibre intake and community composition. The obese microbiota, with its enhanced polysaccharide-degrading capacity, sits at the upper end of this range. The difference between an “efficient” and an “inefficient” community might amount to 100 to 200 kilocalories per day — roughly the energy content of a banana, or a tablespoon of olive oil.

That sounds trivial in the context of a daily intake of 2,000 kilocalories. But weight gain is a cumulative process, not an acute one. An extra 100 kilocalories per day, uncompensated by increased expenditure, translates to roughly five kilograms of fat gain per year. Over a decade, that is the difference between a normal BMI and an obese one. The microbiome's contribution to energy balance is small on any given day but potentially consequential over the timescales on which obesity develops.

There is, however, a paradox here that the chapter must address, because it connects to one of the central themes of the book. SCFAs are, in virtually every other context we have discussed, *beneficial*. Butyrate feeds colonocytes, strengthens the epithelial barrier, induces regulatory T cells, and protects against colorectal cancer. Propionate and acetate have anti-inflammatory properties. A high-fibre diet that promotes SCFA production is, by every measure we have discussed in Chapter 20, a healthy diet. And yet the energy-harvest hypothesis implies that more SCFAs means more calories means more weight gain.

The resolution of this paradox is that SCFAs are not just fuel. They are also signalling molecules. Acetate and propionate activate G-protein-coupled receptors (GPR41 and GPR43) on enteroendocrine L-cells in the gut wall, stimulating the secretion of the satiety hormones GLP-1 and PYY. These hormones slow gastric emptying, reduce appetite, and improve insulin sensitivity. So the same SCFAs that contribute a few extra calories also trigger a hormonal cascade that tends to *reduce* total food intake. In a healthy, diverse community with normal SCFA production, the signalling effect probably outweighs the caloric contribution. In a dysbiotic community — or in a host whose satiety signalling is impaired by other factors — the balance may tip the other way.

This is a recurring lesson of the book: microbial metabolites are not inherently good or bad. They are context-dependent signals whose effects depend on the quantity produced, the receptors available, and the state of the host. Butyrate is beneficial in the colon and essential for barrier integrity; the same butyrate, produced in excess by an overly efficient fermenting community in a host who is already insulin-resistant, contributes to a caloric surplus that exacerbates the problem. Ecology, as we have said before, is not morality.

The bile acid pathway

The second metabolic pathway is bile acid signalling, and it is one we have already explored in some detail. In Chapter 15, we traced the enterohepatic circulation of bile acids, their microbial transformation from primary to secondary forms by gut bacteria (principally *Clostridium* species performing 7 α -dehydroxylation), and their dual role as digestive detergents and metabolic hormones acting through the nuclear receptor FXR and the membrane receptor TGR5.

I will not repeat that biology here — the reader who has come this far will remember it, and the reader who has picked up Volume 2 without Volume 1 can find it summarised in the recap at the front of this book. What matters for the present chapter is what happens to bile acid signalling when the microbiome is disrupted by obesity.

In obese individuals, the gut microbial community produces a different profile of secondary bile acids than in lean individuals. The ratio of primary to secondary bile acids shifts, and the specific secondary species that predominate change. This matters because different bile acids activate FXR and TGR5 with different potencies. In the ileum, FXR activation triggers the secretion of fibroblast growth factor 19 (FGF19), which travels to the liver and suppresses *de novo* lipogenesis — the liver's manufacture of new fat. In the obese microbiome, altered bile acid signalling reduces FGF19

production, releasing the brake on hepatic fat synthesis. The liver makes more fat. That fat accumulates. The result, as we will see in §27.3, is the beginning of a pathway that leads to fatty liver disease.

TGR5, meanwhile, is expressed on enteroendocrine L-cells — the same cells that respond to SCFAs — and its activation by secondary bile acids is one of the triggers for GLP-1 secretion. In an obese microbiome with altered bile acid metabolism, TGR5 signalling is impaired, GLP-1 secretion is reduced, and the host loses one of its main hormonal brakes on appetite and glucose intolerance.

The bile acid story has received dramatic confirmation from bariatric surgery research. Bariatric surgery — gastric bypass and sleeve gastrectomy — is the most effective long-term treatment for severe obesity, producing sustained weight loss of twenty to thirty per cent and, in many patients, complete remission of type 2 diabetes. For decades, the mechanism was assumed to be purely mechanical: a smaller stomach means less food. But the metabolic improvements after surgery are too rapid and too extensive to be explained by caloric restriction alone. Diabetes remission often occurs within days of surgery, before significant weight loss has occurred. Something else is happening.

That something, it is now clear, involves the microbiome and bile acids. A landmark 2022 study by Julia Münzker and colleagues demonstrated that the metabolic benefits of gastric bypass in diet-induced obese rats were largely mediated by the surgery-altered gut microbiota [REF:munzker2022]. When the researchers depleted the gut microbiome with antibiotics after surgery, the metabolic improvements largely disappeared. When they transplanted the post-surgery microbiome into obese recipients who had not been operated on, the recipients showed reduced adiposity, improved glucose homeostasis, and reactivation of thermogenic brown adipose tissue. The mechanism was traced to altered taurine metabolism by the post-surgery community, leading to increased taurine-conjugated secondary bile acids, which activated intestinal FXR and systemic TGR5 to stimulate energy expenditure.

A more recent study confirmed and extended this finding for sleeve gastrectomy. Shaoqian Zhao and colleagues showed that sleeve gastrectomy enriched *Clostridia* species in the gut, increased conjugated secondary bile acids (particularly glycodeoxycholic acid and taurodeoxycholic acid), and that these changes promoted lipolysis and fatty acid oxidation in adipose tissue via TGR5 signalling [REF:zhao2025]. Transplanting post-surgery faeces into obese mice replicated the effect.

The picture that emerges is this: bariatric surgery does not merely shrink the stomach. It restructures the gut microbial community, which restructures bile acid metabolism, which restructures metabolic signalling throughout the body. The surgery is a blunt instrument that happens to produce a precise ecological consequence. Understanding that consequence — and eventually reproducing it without surgery — is one of the most promising avenues in metabolic microbiome research.

The appetite axis and the GLP-1 revolution

The third pathway brings us to the most dramatic intersection of microbiome science and contemporary medicine: the GLP-1 receptor agonists.

Glucagon-like peptide 1 (GLP-1) is a hormone secreted by the enteroendocrine L-cells of the small intestine and colon in response to nutrient ingestion. Its effects include stimulating insulin secretion, suppressing glucagon, slowing gastric emptying, and — crucially for obesity — signalling satiety to the brain. In healthy physiology, the microbiome modulates GLP-1 secretion through both the SCFA and bile acid pathways we have just described: SCFAs activate GPR41/GPR43 on L-cells,

and secondary bile acids activate TGR5 on L-cells, both triggering GLP-1 release. The microbiome is, in this sense, an upstream regulator of one of the body's most important appetite-control signals.

The pharmaceutical exploitation of GLP-1 signalling has produced the most significant advance in obesity medicine in a generation. Semaglutide (marketed as Ozempic for diabetes, Wegovy for weight loss) and the dual GIP/GLP-1 agonist tirzepatide (Mounjaro) achieve sustained weight loss of fifteen to twenty-five per cent in clinical trials — numbers that approach those of bariatric surgery. They have become the fastest-adopted drug class in pharmaceutical history, and they have, quite reasonably, shifted the conversation about obesity from “lifestyle disease requiring willpower” to “treatable medical condition requiring medication.”

What has only recently become apparent is that GLP-1 agonists themselves alter the gut microbiome. A 2025 systematic review of thirty-eight studies found that liraglutide promotes the growth of beneficial genera associated with improved metabolic function, while semaglutide increases *Akkermansia muciniphila* — a species we will return to shortly — but also decreases overall microbial diversity [REF:gofron2025]. The effects vary between drugs, between species studied (mouse versus human), and between study durations, and no consistent “GLP-1 agonist microbiome signature” has yet been established. But the direction is clear: these drugs do not merely mimic a hormone. They reshape the ecosystem that produces it.

One recent study offers a particularly elegant mechanistic vignette. Rodrigo Soares da Silva and colleagues treated high-fat-diet-fed obese mice with semaglutide and found that the drug promoted the growth of acetate-producing bacteria — particularly *Bacteroides acidifaciens* and *Blautia coccoides* — leading to increased acetate levels in the hypothalamus [REF:dasilva2024]. In the hypothalamic arcuate nucleus, acetate activated the GPR43 receptor on POMC neurons (the neurons that suppress appetite) via a PI3K-Akt signalling cascade, contributing to reduced food intake. The drug, in other words, enlisted the microbiome as an intermediary in its own appetite-suppressing mechanism: semaglutide □ altered microbiota □ more acetate □ hypothalamic POMC activation □ less eating.

I find this feedback loop remarkable, and it exemplifies a principle that runs through this entire book: the microbiome is not a bystander in pharmacology. It is a participant. The drug changes the ecosystem, and the changed ecosystem contributes to the drug's effect. Whether this microbial contribution is large enough to matter clinically — whether, for instance, antibiotic-induced disruption of the gut microbiome might blunt the efficacy of semaglutide — is a question that has not yet been answered, but it is one that the field is beginning to ask.

27.3 Type 2 Diabetes, Fatty Liver Disease, and Cardiovascular Risk

The previous chapter dealt with type 1 diabetes — an autoimmune destruction of insulin-producing cells, with a microbial backstory centred on the developing infant immune system. Type 2 diabetes is a different disease. It is not autoimmune. It does not require the immune system to attack the pancreas. Instead, it develops through a slow, progressive impairment of insulin signalling — first in the liver and muscle (insulin resistance), then in the pancreatic beta cells themselves (which eventually fail under the strain of overproduction) — driven by the metabolic consequences of excess adiposity, chronic low-grade inflammation, and, as we will see, the metabolic output of a dysbiotic gut community.

Type 2 diabetes is the metabolic condition in which the microbiome evidence is most advanced, the

mechanistic pathways most clearly mapped, and the therapeutic implications most tangible. It is also the condition in which an important methodological lesson was learned the hard way.

The metformin confound

In 2012, Junjie Qin and colleagues published the first metagenome-wide association study of gut microbiota in type 2 diabetes, using a Chinese cohort of 345 individuals [REF:qin2012]. They found that patients with T2D had decreased butyrate-producing bacteria, increased opportunistic pathogens, and an enrichment of genes involved in oxidative stress resistance and sulphate reduction. It was a substantial paper, and it was widely cited as evidence that T2D has a distinctive microbial signature.

Three years later, Kathrin Forslund and colleagues demonstrated that part of that signature was not due to the disease at all [REF:forslund2015]. It was due to metformin.

Metformin is the first-line drug for type 2 diabetes worldwide, and the vast majority of diagnosed T2D patients take it. Forslund's study, using data from three European cohorts, showed that when you statistically controlled for metformin use, much of the previously reported "T2D microbiome signature" disappeared or weakened. Metformin itself has profound effects on the gut microbiome — it increases *Escherichia*, increases *Akkermansia muciniphila*, alters bile acid metabolism, and shifts SCFA production. We covered these effects in Chapter 22, and they are real and clinically significant. But they mean that any study comparing the microbiomes of treated T2D patients with those of healthy controls is comparing a drugged ecosystem with an undrugged one, and attributing the differences to the disease rather than the drug.

This was a sobering lesson for the entire microbiome field, not just the diabetes subfield. It reinforced a principle we have seen repeatedly: confounders in microbiome studies are everywhere, and the biggest confounders are often the drugs the patients are taking. The Forslund paper did not invalidate the microbiome-T2D connection — there are real microbial differences between pre-diabetic and healthy individuals, observed before any drug treatment — but it demonstrated that the effect is smaller and more nuanced than the early studies suggested, and that disentangling disease from treatment requires careful study design.

The metabolites that matter

With the taxonomic picture complicated by confounders, the field has increasingly focused on microbial *metabolites* rather than microbial *species* as the mediators of the microbiome-diabetes connection. Several metabolites have emerged as particularly important.

The SCFAs and bile acids we discussed in §27.2 are central. Reduced butyrate production, impaired FXR/TGR5 signalling, and the downstream consequences for GLP-1 secretion and hepatic lipogenesis are all part of the T2D picture, and they do not need to be re-explained here. But two additional metabolites deserve specific attention.

The first is **imidazole propionate**. This molecule, produced by gut bacteria from the amino acid histidine, was identified in 2018 as elevated in the portal blood of individuals with type 2 diabetes. Mechanistic studies showed that imidazole propionate impairs insulin signalling by activating the mTORC1 pathway — a central regulator of cell growth and metabolism — in a way that blocks the insulin receptor's downstream cascade [REF:zhu2020metabolites]. It is, in effect, a microbially produced insulin resistance factor: the bacteria make it, the blood carries it, and the liver responds to it by becoming less sensitive to insulin. The finding has not yet translated into a therapeutic

intervention, but it has provided one of the clearest examples of a gut-derived metabolite directly driving a diabetic phenotype.

The second metabolite is **branched-chain amino acids (BCAAs)** — leucine, isoleucine, and valine — whose circulating levels have been associated with insulin resistance and T2D risk in multiple large cohort studies. The gut microbiome is a significant producer of BCAAs, and certain microbial species — notably *Prevotella copri* and *Bacteroides vulgatus* — have been shown to increase circulating BCAA levels when transplanted into germ-free mice. Whether microbial BCAA production is a major driver of insulin resistance in humans, or merely a biomarker of the metabolic state, remains debated. But the association is strong enough to have placed BCAAs alongside TMAO (which we will discuss shortly) as one of the leading candidates for a microbiome-derived cardiometabolic risk factor.

The liver: from NAFLD to MASLD

The metabolic syndrome does not affect the pancreas alone. It affects the liver — and the liver, as we saw in Chapter 26's discussion of primary sclerosing cholangitis, is an organ whose health is intimately connected to the gut microbiome through the portal circulation.

Non-alcoholic fatty liver disease — recently renamed metabolic dysfunction-associated steatotic liver disease, or MASLD, by a 2023 international consensus that sought to connect the condition more explicitly to its metabolic roots — is the hepatic manifestation of the metabolic syndrome. It ranges from simple steatosis (fat accumulation in liver cells) through steatohepatitis (inflammation and cell damage) to fibrosis, cirrhosis, and hepatocellular carcinoma. It affects roughly a quarter of the world's adult population. It is, by prevalence, the most common liver disease on earth, and it has no approved pharmacological treatment — although resmetirom, a thyroid hormone receptor agonist, received approval in 2024 for the steatohepatitis stage, and semaglutide is in advanced trials.

The microbiome's role in MASLD is mediated through the same gut-liver axis we have discussed throughout this book, but with specific features that deserve attention.

The mechanism is, in outline, familiar: gut dysbiosis leads to increased intestinal permeability, translocation of bacterial products (particularly LPS) into the portal circulation, activation of hepatic Toll-like receptor 4 (TLR4) on Kupffer cells (the liver's resident macrophages), and a chronic inflammatory response that promotes fat accumulation, cell damage, and fibrosis. A 2024 review by Raquel Benedé-Ubieto and colleagues in *Gut Microbes* provides the most comprehensive synthesis of this pathway, framing MASLD as a disorder of the gut-liver barrier [REF:benedeubieto2024].

The microbial signatures of MASLD are, characteristically, a mirror image of those seen in metabolic health: decreased *Akkermansia muciniphila* (the mucus-layer specialist that strengthens barrier integrity), decreased *Faecalibacterium* (the main butyrate producer), and increased *Escherichia* and *Prevotella* [REF:long2024]. These shifts are associated with increased gut permeability, increased serum LPS, and increased hepatic inflammation — the same downstream cascade.

But the MASLD story adds a specific metabolic dimension that goes beyond inflammation. Gut bacteria alter the bile acid pool (as we discussed in Chapter 15 and §27.2), and altered bile acid signalling through FXR promotes hepatic *de novo* lipogenesis — the liver's production of fat from carbohydrate precursors. They also metabolise dietary choline, reducing its availability for the synthesis of phosphatidylcholine — a molecule the liver needs to package and export fat as VLDL particles. When choline is depleted by microbial metabolism, the liver cannot export fat efficiently,

and the fat accumulates. This choline-depletion pathway was one of the earliest identified mechanisms linking gut microbes to fatty liver, and it remains one of the most mechanistically convincing.

The heart: TMAO and the microbiome-cardiovascular connection

If the microbiome-liver connection runs through bile acids and choline, the microbiome-heart connection runs through a molecule with a name that has become, in the popular press, almost as famous as cholesterol: trimethylamine N-oxide, or TMAO.

The TMAO story began in 2011, when a team led by Stanley Hazen at the Cleveland Clinic published a paper in *Nature* that fundamentally changed how cardiologists think about dietary risk [REF:wang2011]. Using a metabolomics screen — an unbiased survey of all small molecules in the blood — they identified three metabolites whose levels predicted cardiovascular events: choline, betaine, and TMAO. All three are linked by a common pathway. Dietary choline (found in eggs, liver, and other animal products) and L-carnitine (found in red meat) are metabolised by gut bacteria to produce trimethylamine (TMA). TMA is absorbed into the portal blood, transported to the liver, and oxidised by the enzyme flavin monooxygenase 3 (FMO3) to produce TMAO.

TMAO, the Hazen group showed, promotes atherosclerosis through multiple mechanisms: it enhances foam cell formation (the cholesterol-laden macrophages that build arterial plaque), increases platelet hyperreactivity (promoting blood clot formation), induces vascular inflammation, and impairs reverse cholesterol transport (the process by which cholesterol is removed from arteries). In animal models, supplementing the diet with TMAO or its precursors accelerated atherosclerosis, and germ-free mice — which cannot produce TMA because they lack gut bacteria — were protected. The pathway was explicitly microbial: no gut microbes, no TMA, no TMAO, less atherosclerosis.

Subsequent epidemiological studies confirmed that circulating TMAO levels predict cardiovascular events — heart attacks, strokes, and cardiovascular death — independently of traditional risk factors like cholesterol, blood pressure, and smoking. A comprehensive review by Marco Witkowski, Taylor Weeks, and Hazen himself summarises the evidence and the mechanisms, and it is the standard reference for the field [REF:witkowski2020]. The same group has since identified a second microbiome-derived cardiovascular metabolite, **phenylacetylglutamine (PAGln)**, produced by gut bacterial metabolism of phenylalanine, which promotes adverse cardiovascular phenotypes through interaction with adrenergic receptors.

The TMAO story is one of the most elegant in microbiome science — a clean pathway from dietary substrate through microbial metabolism to host disease — and it has been enormously influential. But I want to add a note of caution, because the field deserves it.

In 2023, Marina Canyelles and colleagues published a critical review asking whether TMAO is truly a *causal* factor in cardiovascular disease, or whether it is primarily a *marker* of something else — specifically, of impaired renal function [REF:canyelles2023]. TMAO is eliminated by the kidneys, and its circulating levels rise as kidney function declines. When epidemiological studies adjust for renal function, the association between TMAO and cardiovascular outcomes weakens or disappears in many cohorts. A Mendelian randomisation study — a genetic approach that tests causality by using naturally occurring genetic variants as instruments — did not find a significant association between genetically predicted higher TMAO levels and cardiometabolic disease. A second Mendelian randomisation study found a positive causal relationship between TMAO and systolic blood pressure, which could partly explain the cardiovascular link through a hypertension-and-

kidney-function pathway rather than through direct atherogenic effects.

This does not mean that TMAO is irrelevant. The animal model evidence for direct atherogenic effects is strong, and the Mendelian randomisation analyses have limitations (they capture lifelong genetically determined TMAO levels, which may not reflect the dynamic changes caused by dietary shifts). But it does mean that the popular narrative — “red meat □ gut bacteria □ TMAO □ heart attack” — is more complicated than it sounds, and that the contribution of TMAO to cardiovascular disease in humans may be partially or substantially mediated by its relationship with kidney function rather than by a direct effect on arterial walls.

I raise this not to dismiss the TMAO story but to illustrate a principle that applies throughout this chapter and this book: the gap between a compelling mechanistic pathway and a validated clinical target is real, and the microbiome field has repeatedly stumbled in that gap. TMAO may eventually prove to be a druggable target for cardiovascular prevention. It may also turn out to be a useful biomarker that reflects disease risk without being a primary driver. The evidence, as of 2026, supports both interpretations, and the honest answer is that we do not yet know which is correct.

27.4 *Akkermansia muciniphila* and the Next-Generation Probiotic Hope

If there is a single microbial species that has become the poster organism for metabolic microbiome research, it is *Akkermansia muciniphila* — the mucus-dwelling bacterium we first met in Chapter 5 as a specialist of the colonic mucus layer, and which has since appeared in virtually every chapter of this book where the words “barrier integrity” or “metabolic health” are mentioned.

The case for *Akkermansia* in metabolic disease rests on three pillars. First, its abundance is consistently reduced in obesity, type 2 diabetes, and MASLD — across dozens of studies and multiple populations. Second, its mechanism of action is plausible: by feeding on mucus glycans and stimulating goblet cells to produce more mucus, it strengthens the gut barrier, reduces LPS translocation, and attenuates the chronic low-grade inflammation that drives metabolic dysfunction. And third, it has been tested in humans.

The human trial that matters most was published in *Nature Medicine* in 2019 by Clara Depommier, Patrice Cani, and colleagues at UCLouvain in Belgium [REF:depommier2019]. It was a randomised, double-blind, placebo-controlled study in overweight and obese insulin-resistant volunteers — forty enrolled, thirty-two completed. One group received live *A. muciniphila*, another received pasteurised (heat-killed) *A. muciniphila*, and a third received placebo, all for three months.

The results were striking, and they were also surprising. The pasteurised form outperformed the live form. Compared with placebo, pasteurised *A. muciniphila* improved insulin sensitivity by about twenty-nine per cent, reduced fasting insulin by thirty-four per cent, and lowered total cholesterol by about nine per cent. It also slightly reduced body weight (by about 2.3 kilograms) and hip circumference, although these changes did not quite reach statistical significance. The live form, unexpectedly, showed smaller and generally non-significant effects.

This was a proof-of-concept study — small, short, and designed to test safety and tolerability rather than to prove clinical efficacy. It proved both. But it also raised as many questions as it answered. Why did the dead form work better than the live one? The leading hypothesis is that a specific protein on the outer membrane of *A. muciniphila* — Amuc_1100, a pilus-like protein — is the active ingredient, and that pasteurisation exposes it in a form that is more accessible to the host’s immune

receptors (particularly TLR2) than the intact living bacterium. If that is true, then the therapeutic agent is not a probiotic organism but a postbiotic molecule — a structural component of a dead bacterium, acting as a drug rather than as a living ecosystem member.

Since Depommier, several further trials have explored *Akkermansia* in metabolic contexts, including a 2024 study showing that oral *A. muciniphila* augmented glycaemic control in type 2 diabetes patients alongside standard care [REF:attaye2024]. The results are consistently positive in direction but consistently modest in magnitude, and no study has yet demonstrated the kind of clinically transformative weight loss that patients and investors are hoping for.

Akkermansia is, in my view, the honest emblem of the metabolic microbiome field: a genuinely promising organism with a plausible mechanism, real but modest human evidence, and a distance still to travel before it can be called a treatment rather than a hope. The next-generation probiotic company that has commercialised pasteurised *A. muciniphila* as a supplement is selling, for the moment, ahead of the evidence — a situation that should sound familiar from Chapter 23's discussion of probiotics more broadly.

27.5 Why Microbiome-Based Weight Loss Therapies Haven't Arrived Yet

It is now twenty years since Ruth Ley and Jeffrey Gordon published the founding papers of the obesity-microbiome field. In that time, the science has advanced enormously. We understand, in molecular detail, how microbial metabolites influence energy harvest, bile acid signalling, appetite regulation, insulin sensitivity, hepatic lipogenesis, and cardiovascular risk. We have identified specific species and specific metabolites that correlate with metabolic health and metabolic disease. We have shown, in animal models, that transplanting the microbiome can transfer metabolic phenotypes.

And yet no microbiome-based therapy for obesity has reached clinical practice. No probiotic produces meaningful weight loss. No prebiotic reverses type 2 diabetes. Faecal microbiota transplantation from lean donors temporarily improves insulin sensitivity in metabolic syndrome but does not produce sustained weight loss [REF:vrieze2012]. The gap between the science and the clinic remains wide.

Why?

The first reason is the oldest one in biology: **correlation is not causation**, and proving causation in humans is hard. The animal models are compelling — germ-free mice colonised with obese microbiota gain fat, bariatric surgery microbiome transplants improve metabolism in recipients — but mice are not people. Human obesity involves decades of interaction between genetics, diet, activity, stress, sleep, medication, and microbial ecology. Isolating the microbial contribution from that tangle requires a kind of experimental control that is, for obvious ethical reasons, impossible in humans.

The second reason is **individual variation**. We have known since the Zeevi 2015 personalised nutrition study — discussed in Chapter 20 — that glycaemic responses to identical foods differ enormously between individuals, and that a substantial fraction of that variation is explained by the microbiome. This means that a microbiome intervention that works in one person may not work in another, not because the intervention is wrong but because their ecosystems are different. A one-size-fits-all probiotic for weight loss is, on first principles, unlikely to succeed. The field needs

personalised approaches, and personalised approaches are harder to develop, harder to test, and harder to regulate.

The third reason is **ecological complexity**. The gut microbiome is not a single lever that can be pulled. It is a network of several hundred species, interacting with each other and with the host, producing thousands of metabolites, and responding dynamically to every meal, every medication, and every perturbation. Adding a single probiotic species to this network is like introducing one new employee to a company of five hundred: the existing employees will either absorb the newcomer into the prevailing culture or reject it outright. The colonisation resistance that protects a healthy microbiome from pathogens (Chapter 5) also protects it from probiotics. Changing the ecosystem requires changing the *ecology*, not adding a species.

The fourth reason is that the bar has moved. In 2005, when the field began, the best available obesity pharmacotherapy produced about five per cent weight loss — barely clinically relevant. A microbiome-based approach that achieved even modest improvements in body weight or insulin sensitivity would have been competitive. In 2026, semaglutide and tirzepatide achieve fifteen to twenty-five per cent weight loss. The GLP-1 revolution has set a therapeutic standard that no microbiome intervention is remotely close to matching. The question is no longer “can the microbiome treat obesity?” but “can the microbiome *add something* to treatments that are already transformative?”

I think the answer to that reframed question is yes, but not in the way the popular press has imagined. The microbiome’s greatest near-term contribution to metabolic medicine is unlikely to be a standalone therapy. It is more likely to be threefold.

First, **pharmacomicrobiomics**: understanding how the microbiome modulates the response to existing drugs. Metformin’s efficacy is partly microbial (Chapter 22). GLP-1 agonists reshape the microbiome in ways that may contribute to their effects. Identifying which patients will respond best to which drugs, based on their microbial profiles, is a realistic near-term goal.

Second, **risk stratification**: using microbial gene richness, specific metabolite levels (TMAO, imidazole propionate, BCAAs), or community composition to identify individuals at high risk of metabolic disease progression *before* they develop overt diabetes or cardiovascular events. The Le Chatelier finding — that low microbial gene count predicts future weight gain — is a proof of concept for this approach.

Third, **adjunctive ecological interventions**: dietary strategies (fibre type and quantity, fermented foods, chrononutrition) that are tailored to an individual’s microbiome to optimise metabolic outcomes alongside pharmacotherapy. This is not a probiotic capsule. It is a personalised dietary prescription informed by microbial profiling — the precision nutrition that Chapter 20 began to sketch and that this chapter’s science makes increasingly feasible.

None of these is as dramatic as “a pill that cures obesity.” But all of them are within reach, and all of them are more likely to improve patient outcomes in the next decade than any single-species probiotic or microbiome transplant protocol.

The honesty the field needs

I want to end this chapter where I ended the last one: with a plea for honesty about what the science does and does not support, because the stakes are high and the misinformation is abundant.

The microbiome is part of the metabolic disease story. It is not the main character and it is not

the solution. Anyone who tells you that a probiotic can cure your diabetes, that a particular fermented food will melt your belly fat, or that your microbiome test reveals the “root cause” of your metabolic syndrome is selling something that the science does not support. The mechanisms are real. The clinical applications are not yet here. The distance between the two is measured not in marketing but in randomised controlled trials, and those trials have, so far, delivered results that are interesting but not transformative.

What the microbiome model does offer — and this is not nothing — is a richer understanding of why metabolic diseases behave the way they do. It explains why identical twins can have different metabolic fates. It explains why weight loss after bariatric surgery is more durable than weight loss from dieting. It explains why metformin works better in some patients than others. It explains why the calorie is not, after all, a sufficient unit of nutritional accounting — because the same calorie, processed by a different microbial community, has different metabolic consequences.

That understanding is worth having. It changes how we think about the diseases that kill more people, worldwide, than any other. And it will, eventually, change how we treat them. The treatments will be ecological, personalised, and modest in their ambitions — not a revolution but a refinement. Given the scale of the problem, even a refinement would save millions of lives.

Key Chapter References

- [REF:ley2005] Ley RE, Bäckhed F, Turnbaugh P, et al. (2005). Obesity alters gut microbial ecology. *Proc Natl Acad Sci USA*. 102(31):11070–5. doi:10.1073/pnas.0504978102
- [REF:ley2006] Ley RE, Turnbaugh PJ, Klein S, Gordon JI (2006). Human gut microbes associated with obesity. *Nature*. 444(7122):1022–3. doi:10.1038/4441022a
- [REF:turnbaugh2006] Turnbaugh PJ, Ley RE, Mahowald MA, et al. (2006). An obesity-associated gut microbiome with increased capacity for energy harvest. *Nature*. 444(7122):1027–31. doi:10.1038/nature05414
- [REF:lechatelier2013] Le Chatelier E, Nielsen T, Qin J, et al. (2013). Richness of human gut microbiome correlates with metabolic markers. *Nature*. 500(7464):541–6. doi:10.1038/nature12506
- [REF:qin2012] Qin J, Li Y, Cai Z, et al. (2012). A metagenome-wide association study of gut microbiota in type 2 diabetes. *Nature*. 490(7418):55–60. doi:10.1038/nature11450
- [REF:forslund2015] Forslund K, Hildebrand F, Nielsen T, et al. (2015). Disentangling type 2 diabetes and metformin treatment signatures in the human gut microbiota. *Nature*. 528(7581):262–266. doi:10.1038/nature15766
- [REF:wang2011] Wang Z, Klipfell E, Bennett BJ, et al. (2011). Gut flora metabolism of phosphatidylcholine promotes cardiovascular disease. *Nature*. 472(7341):57–63. doi:10.1038/nature09922
- [REF:witkowski2020] Witkowski M, Weeks TL, Hazen SL (2020). Gut Microbiota and Cardiovascular Disease. *Circ Res*. 127(4):553–570. doi:10.1161/CIRCRESAHA.120.316242
- [REF:canyelles2023] Canyelles M, Borràs C, Rotllan N, Tondo M, Escolà-Gil JC, Blanco-Vaca F (2023). Gut Microbiota-Derived TMAO: A Causal Factor Promoting Atherosclerotic Cardiovascular Disease? *Int J Mol Sci*. 24(3):1940. doi:10.3390/ijms24031940

- [REF:depommier2019] Depommier C, Everard A, Druart C, et al. (2019). Supplementation with *Akkermansia muciniphila* in overweight and obese human volunteers: a proof-of-concept exploratory study. *Nat Med*. 25(7):1096–1103. doi:10.1038/s41591-019-0495-2
- [REF:attaye2024] Attaye I, et al. (2024). Oral *Akkermansia muciniphila* L2-7 augments glycemic control in type 2 diabetes. *iScience*. 27(8):110455. doi:10.1016/j.isci.2024.110455
- [REF:masi2026] Masi D, Watanabe M, Clément K (2026). Gut microbiome and obesity care: bridging dietary, surgical, and pharmacological interventions. *Cell Rep Med*. 7(2):102573. doi:10.1016/j.xcrm.2025.102573
- [REF:munzker2022] Münzker J, Haase N, Till A, et al. (2022). Functional changes of the gastric bypass microbiota reactivate thermogenic adipose tissue and systemic glucose control via intestinal FXR-TGR5 crosstalk in diet-induced obesity. *Microbiome*. 10(1):96. doi:10.1186/s40168-022-01264-5
- [REF:zhao2025] Zhao S, Lin H, Li W, et al. (2025). Post sleeve gastrectomy-enriched gut commensal *Clostridia* promotes secondary bile acid increase and weight loss. *Gut Microbes*. 17(1):2462261. doi:10.1080/19490976.2025.2462261
- [REF:gofron2025] Gofron KK, Wasilewski A, Małgorzewicz S (2025). Effects of GLP-1 Analogues and Agonists on the Gut Microbiota: A Systematic Review. *Nutrients*. 17(8):1303. doi:10.3390/nu17081303
- [REF:dasilva2024] da Silva RS, de Paiva IHR, Mendonça IP, et al. (2024). Anorexigenic and anti-inflammatory signaling pathways of semaglutide via the microbiota-gut-brain axis in obese mice. *Inflammopharmacology*. 33(2):845–864. doi:10.1007/s10787-024-01603-y
- [REF:benedeubieto2024] Benedé-Ubieto R, Cubero FJ, Nevzorova YA (2024). Breaking the barriers: the role of gut homeostasis in Metabolic-Associated Steatotic Liver Disease (MASLD). *Gut Microbes*. 16(1):2331460. doi:10.1080/19490976.2024.2331460
- [REF:long2024] Long Q, Luo F, Li B, et al. (2024). Gut microbiota and metabolic biomarkers in metabolic dysfunction-associated steatotic liver disease. *Hepatol Commun*. 8(3). doi:10.1097/HC9.0000000000000310
- [REF:vrieze2012] Vrieze A, Van Nood E, Holleman F, et al. (2012). Transfer of intestinal microbiota from lean donors increases insulin sensitivity in individuals with metabolic syndrome. *Gastroenterology*. 143(4):913–916.e7. doi:10.1053/j.gastro.2012.06.031
- [REF:zhu2020metabolites] Zhu T, Goodarzi MO (2020). Metabolites Linking the Gut Microbiome with Risk for Type 2 Diabetes. *Curr Nutr Rep*. 9(2):83–93. doi:10.1007/s13668-020-00307-3

Chapter 28: Cancer

28.0 The Oldest Suspicion

The idea that infection causes cancer is not new. It is, in fact, one of the oldest ideas in oncology — older than oncology itself, if we count the observations of eighteenth-century physicians who noticed that cancers sometimes arose in tissues scarred by chronic infection. But for most of the twentieth century, the idea was marginal. Cancer was understood as a disease of the cell: mutations, chromosomal rearrangements, the unleashing of oncogenes and the silencing of tumour suppressors. The external agents that mattered were radiation and chemical carcinogens — mutagens that damaged DNA directly. Microbes, with a few exotic exceptions such as the hepatitis B virus and gastric cancer's eventual villain, were not part of the mainstream story.

That has changed. We now know that approximately one in five human cancers worldwide — perhaps twenty per cent of the global burden — is attributable to infectious agents. Most of the established links are viral: human papillomavirus and cervical cancer, Epstein-Barr virus and several lymphomas, hepatitis B and C viruses and hepatocellular carcinoma. But the microbiome revolution has added a new dimension to this picture, one that goes beyond the classical “one pathogen, one cancer” model. The bacteria that live in and on us — the commensal communities we have spent this book exploring — are now implicated in cancer in at least four distinct ways.

First, specific bacterial species can directly promote the development of tumours, either by damaging DNA, by hijacking host cell signalling, or by creating a chronically inflamed tissue environment in which mutations accumulate faster than they can be repaired. *Helicobacter pylori* and gastric cancer is the paradigm; *Fusobacterium nucleatum* and colorectal cancer is the rising story.

Second, the broader composition of the gut microbiome — its diversity, its metabolic output, its tonic influence on the immune system — determines how effectively the body fights cancers once they arise. This is the immunotherapy story we introduced in Chapter 12, and we will deepen it here.

Third, the microbiome is emerging as a predictive tool — a biological readout that can tell clinicians, before treatment begins, how likely a cancer is to recur or to respond to a given therapy. The microbial fingerprint study we met in Chapter 12 is the most recent example of this, and it points toward a future in which stool samples sit alongside tumour biopsies in the oncologist's diagnostic kit.

Fourth, and most speculatively, the microbiome may offer a route to early cancer detection — a non-invasive screen based on the microbial shifts that accompany early-stage tumours, potentially years before symptoms appear.

This chapter takes these four threads in order. It is not a comprehensive catalogue of every cancer in which the microbiome has been implicated — that list grows longer each year and would fill a textbook of its own. Instead, it focuses on the cases in which the evidence is strongest, the mechanisms are best understood, and the clinical implications are most tangible. The reader who finishes this chapter should come away with a clear picture of what the microbiome can and cannot yet tell us about cancer, and why the field believes the next decade will transform that picture dramatically.

28.1 *Helicobacter pylori* and Gastric Cancer — The Original Microbe-Cancer Link

The bacterium that rewrote medicine

In the early 1980s, two Australian researchers at the Royal Perth Hospital did something that medical orthodoxy said was unnecessary: they looked carefully at the bacteria growing in the stomachs of patients with gastritis and peptic ulcers. Barry Marshall, a trainee in internal medicine, and Robin Warren, a pathologist, had noticed curved bacilli colonising the gastric mucosa of biopsy specimens — organisms that the conventional wisdom said could not exist there. The stomach, with its bath of hydrochloric acid at pH 1–2, was supposed to be sterile. Any bacterium that found its way in would be dissolved. The textbooks said so.

The textbooks were wrong. The organism — eventually named *Helicobacter pylori* — had evolved an elegant survival strategy. It produces copious amounts of urease, an enzyme that splits urea into ammonia and carbon dioxide. The ammonia neutralises the acid in the bacterium's immediate vicinity, creating a microenvironment of near-neutral pH in which *H. pylori* can thrive, nestled in the mucus layer that lines the gastric epithelium. It does not merely tolerate the stomach; it has made the stomach its exclusive ecological niche. No other known bacterium has colonised this environment with comparable success.

Marshall and Warren's claim that this organism caused gastritis and peptic ulcers — conditions that the medical establishment attributed to stress, diet, and excess acid production — was met with fierce resistance. The story of Marshall's self-experimentation in 1984, when he drank a broth culture of *H. pylori* to prove it caused gastritis (it did, spectacularly), has become one of the most famous anecdotes in the history of medicine. It took another decade for the evidence to become overwhelming enough to change clinical practice. By the mid-1990s, antibiotic-based eradication therapy had replaced antacids and surgical vagotomy as the standard treatment for peptic ulcer disease. In 2005, Marshall and Warren received the Nobel Prize in Physiology or Medicine.

But the implications went far beyond ulcers. In 1994 — before Marshall and Warren had their Nobel — the World Health Organization's International Agency for Research on Cancer classified *H. pylori* as a Group 1 carcinogen: a definitive cause of cancer in humans. It was the first bacterium ever to receive this classification.

The carcinogenic cascade

H. pylori infects approximately half the world's population, with prevalence ranging from around thirty per cent in affluent Western countries to over eighty per cent in parts of sub-Saharan Africa, South America, and South Asia. Yet only one to three per cent of infected individuals develop gastric cancer. The discrepancy tells us that *H. pylori* is necessary but not sufficient — a risk factor, not a guarantee. The outcome depends on three interacting variables: the virulence of the particular strain, the genetic susceptibility of the host, and the environmental context, particularly diet.

The virulence story centres on two proteins: **CagA** and **VacA**.

CagA is the more consequential. It is injected directly into gastric epithelial cells by a molecular syringe called the **type IV secretion system** — a needle-like apparatus that punches through the host cell membrane and delivers the bacterial protein into the cytoplasm. Once inside, CagA

is phosphorylated by host cell kinases and begins to interfere with a remarkable range of intracellular signalling pathways: the ERK/MAPK pathway (cell proliferation), PI3K/Akt (cell survival), NF- κ B (inflammation), Wnt/ β -catenin (cell fate and polarity), and the Hippo pathway (organ size and growth control). The cumulative effect is to push the infected epithelial cell toward a state that resembles the early stages of malignant transformation — increased proliferation, resistance to apoptosis, loss of cell polarity, disruption of cell-cell junctions, and chronic activation of inflammatory signalling [REF:hatakeyama2019].

CagA has been called the only bacterially derived oncoprotein — a bacterial molecule that functions, inside the host cell, much as a viral oncoprotein does. The analogy with the oncoproteins of human papillomavirus (E6 and E7) or Epstein-Barr virus (LMP1) is not perfect, but the functional parallels are striking enough to have reshaped how microbiologists think about bacterial pathogenesis.

VacA — the vacuolating cytotoxin — works differently. Rather than hijacking signalling pathways, it disrupts cellular housekeeping. It forms pores in cell membranes, induces the formation of large intracellular vacuoles, triggers mitochondrial dysfunction, and — critically — impairs autophagy, the cellular recycling machinery that normally clears damaged proteins and organelles. Recent work has shown that VacA and CagA cooperate: in the presence of VacA, CagA accumulates in dysfunctional autophagosomes, amplifying its oncogenic effects because the cell can no longer clear the foreign protein [REF:raju2019].

The long road from infection to cancer

Not all *H. pylori* strains carry CagA, and among those that do, the virulence of the CagA protein varies geographically. East Asian strains carry a CagA variant with a different phosphorylation motif (EPIYA-D rather than EPIYA-C) that binds more avidly to host signalling proteins and induces stronger oncogenic signalling. This is part of the reason why gastric cancer rates are dramatically higher in Japan, South Korea, and China than in Western countries, even after controlling for other risk factors.

But even with the most virulent strain, cancer takes decades. The typical progression — known as the **Correa cascade** after the Colombian pathologist Pelayo Correa, who first described it — moves through a series of histological stages: chronic active gastritis, atrophic gastritis (loss of acid-secreting glands), intestinal metaplasia (replacement of the gastric lining with an intestinal-type epithelium), dysplasia, and finally adenocarcinoma. Each step takes years. The entire cascade may span three to four decades from initial infection to clinical cancer. Along the way, there are opportunities for prevention: eradication of *H. pylori* before atrophic gastritis develops can reduce gastric cancer risk by thirty to forty per cent. After atrophic gastritis is established, eradication still helps, but the risk reduction is smaller — the tissue has already been damaged, and some of the pre-cancerous changes may be self-sustaining even without the bacterium [REF:choi2018].

This is what makes *H. pylori* carcinogenesis a “hit-and-run” phenomenon, at least in its later stages. The bacterium initiates the cascade, sustains it through chronic inflammation and direct oncogenic signalling, and then — in many cases — disappears from the cancerous tissue itself. Established gastric cancers are often *H. pylori*-negative. The bacterium is no longer needed; the genetic and epigenetic damage it inflicted has become self-perpetuating. This pattern — initiate, sustain, and then become dispensable — will reappear, in a different form, when we turn to *Fusobacterium nucleatum* and colorectal cancer.

Beyond *H. pylori*: the gastric microbiome

There is one more layer to the gastric cancer story that the *H. pylori*-centric narrative tends to obscure. *H. pylori* is not the only organism in the stomach. When it colonises the gastric mucosa and raises the local pH through its urease activity, it inadvertently opens the door to other bacteria that cannot normally survive in the acidic environment. The result is a shift in the entire gastric microbiome — a dysbiosis driven by the dominant organism's ecological engineering.

Recent studies have shown that the non-*H. pylori* gastric microbiome differs between patients with simple gastritis, those with pre-cancerous intestinal metaplasia, and those with established gastric cancer. The diversity of the gastric community changes as the tissue progresses along the Correa cascade, and some of the non-*H. pylori* species — including members of the genera *Streptococcus*, *Prevotella*, and *Veillonella* — may contribute to the inflammatory and genotoxic environment independently. Whether these organisms are fellow travellers or co-conspirators is not yet clear. But the emerging picture is that gastric cancer is not solely a disease of one bacterium acting on one host. It is, like so much else in this book, a disease of an ecosystem [REF:castaño-rodríguez2017].

Eradication therapy — typically a combination of a proton pump inhibitor and two antibiotics, taken for ten to fourteen days — has dramatically reduced the incidence of gastric cancer in countries where population-level screening and treatment have been implemented, most notably Japan and South Korea. But eradication is itself a microbiome perturbation. A 2025 systematic review found that triple or quadruple therapy for *H. pylori* significantly reshapes the gut microbiome beyond the stomach, reducing diversity and, in some cases, promoting the expansion of antibiotic-resistant organisms in the intestinal community [REF:eradication_microbiome2025]. The irony is acute: treating the bacterium that causes gastric cancer may, at the margins, contribute to the kind of intestinal dysbiosis that this book has linked to a range of other diseases. It is a trade-off, and on balance the trade-off overwhelmingly favours eradication in patients at significant cancer risk. But it is a reminder that no intervention in a microbial ecosystem comes without unintended consequences.

28.2 *Fusobacterium nucleatum* and Colorectal Cancer

An unexpected intruder

If *H. pylori* is the familiar villain of the microbe-cancer story — discovered in the 1980s, Nobel-recognised, taught in every medical school — then *Fusobacterium nucleatum* is the newcomer that nobody saw coming.

F. nucleatum is an anaerobic, gram-negative, spindle-shaped bacterium whose natural home is the human mouth. It is one of the most abundant organisms in the subgingival plaque that forms below the gumline, and it has been studied for decades in the context of periodontal disease — a world away from oncology. Nobody in the colorectal cancer field was looking for it.

In 2011 and 2012, two independent genomic studies changed that. Castellarin and colleagues, using RNA sequencing of tumour tissue followed by subtraction of human sequences, found that *F. nucleatum* was markedly over-represented in colorectal carcinomas compared with matched normal tissue from the same patients [REF:castellarin2012]. Simultaneously, Kostic and colleagues used whole-genome sequencing of tumour DNA with a similar host-subtraction approach and found an enrichment of *Fusobacterium* sequences in most of the colorectal cancers they examined [REF:kostic2012]. Both groups had stumbled on the same finding using different methods, in

different cohorts, on different continents. The convergence was compelling.

In the years since, *F. nucleatum* has become the most intensively studied bacterium in colorectal cancer after *H. pylori*. The evidence linking it to CRC has progressed from “interesting association” to “mechanistically plausible contributor” — though, as we will see, the question of whether it truly *causes* cancer or merely *accelerates* existing tumours remains a matter of active debate.

How does a mouth bacterium end up in a colon tumour?

The first puzzle is anatomical. *F. nucleatum* is an oral organism. How does it reach the colon, and why does it specifically colonise tumour tissue?

The route appears to be haematogenous — via the bloodstream. *F. nucleatum* enters the circulation through damaged or inflamed gum tissue (periodontal disease is remarkably common, affecting nearly half of adults over thirty), travels through the vasculature, and seeds distant tissues. This is the same route by which oral bacteria are thought to contribute to endocarditis and to the cardiovascular disease associations we discussed in Chapter 9. The colon tumour, however, provides something that healthy colonic tissue does not: a molecular welcome mat.

The key molecule is **Gal-GalNAc** (galactose-*N*-acetylgalactosamine), a disaccharide that is overexpressed on the surface of colorectal cancer cells. *F. nucleatum* carries an outer-membrane protein called **Fap2** that binds specifically to Gal-GalNAc. This is tumour-selective targeting: healthy colon cells express little Gal-GalNAc, so the bacterium passes them by and homes in on the malignant tissue [REF:abed2016]. In 2025, structural work by several groups revealed the molecular details of this interaction, showing how the Fap2 protein recognises and docks onto its target sugar with high specificity — an interaction precise enough that researchers have begun exploring whether it could be exploited therapeutically, using engineered Fap2-based targeting molecules to deliver drugs specifically to CRC tumours [REF:fap2_structure2025].

The mechanistic toolkit

Once *F. nucleatum* has colonised the tumour, it deploys a second set of virulence factors that promote tumour growth and metastasis through at least three mechanisms.

Direct oncogenic signalling. A second adhesin, **FadA**, binds to E-cadherin on the surface of epithelial cells. E-cadherin is a cell-cell junction protein that normally acts as a brake on cell proliferation — when it is engaged in its normal function, holding cells together, it suppresses the Wnt/ β -catenin pathway that drives cell division. When FadA binds E-cadherin, it disrupts this brake, activating β -catenin signalling and promoting proliferation. This is the same pathway that is mutationally activated in the majority of colorectal cancers; *F. nucleatum* gives it an additional push from the outside [REF:rubinstein2013].

Immune evasion. The Fap2 protein has a second, entirely separate function beyond tumour homing. It binds to **TIGIT** (T cell immunoreceptor with Ig and ITIM domains), an inhibitory receptor expressed on natural killer (NK) cells and certain T cell subsets. When Fap2 engages TIGIT, it sends an inhibitory signal that suppresses the cytotoxic activity of these immune cells — effectively telling the immune system to stand down in the vicinity of the tumour. This is immunosuppression by microbial proxy: the bacterium shields the tumour from immune attack, not for the tumour’s benefit (the bacterium has no “interest” in the tumour’s survival), but as a side effect of its own immune evasion strategy [REF:gur2015].

Pro-inflammatory environment. *F. nucleatum* infection in the tumour microenvironment stimulates the recruitment of myeloid-derived suppressor cells (MDSCs), promotes the production of pro-inflammatory cytokines including IL-8 and TNF- α , and creates a local inflammatory milieu that favours tumour progression. Chronic inflammation, as we have discussed throughout this book, is a potent promoter of carcinogenesis — it generates reactive oxygen species that damage DNA, selects for cells with mutations that confer survival advantages in a hostile environment, and provides growth factors and angiogenic signals that help tumours expand.

The clade question

For years, a puzzling inconsistency nagged at the *F. nucleatum*-CRC field. Not every study could replicate the association. Some cohorts showed strong enrichment; others showed weaker signals. Part of the explanation turned out to be that researchers were treating *F. nucleatum* as a single entity when it is, in fact, multiple distinct organisms hiding under one name.

In 2024, a landmark study in *Nature* by Susan Bullman's group at the Fred Hutchinson Cancer Center resolved this. By generating complete, closed genomes for 135 *F. nucleatum* strains — 80 from the mouths of cancer-free individuals and 55 cultured directly from CRC tumours — they showed that the subspecies *F. nucleatum* subsp. *animalis* (Fna), previously considered a single lineage, is in fact composed of two genetically distinct clades: **Fna C1** and **Fna C2** [REF:bullman2024].

Only Fna C2 dominated the colorectal cancer niche. Fna C1 remained largely confined to the oral cavity. The two clades differed in morphology (C2 cells were longer and thinner), in their repertoire of virulence factors (C2 carried additional copies of *fap2* and several other CRC-enriched genes), and in functional consequence: in mouse models, Fna C2 treatment led to increased numbers of intestinal adenomas and distinct shifts in the metabolic profile of the tumour environment.

This finding matters for two reasons. First, it explains much of the previous inconsistency in the literature: studies that failed to distinguish C1 from C2 were diluting the cancer signal with a non-pathogenic relative. Second, it reinforces a theme that runs through this entire book: the species level is often the wrong resolution for understanding the microbiome. What matters — in disease, in health, in ecological function — is frequently the strain, the subspecies, or in this case the clade. Treating “*Fusobacterium nucleatum*” as a single actor is like treating “dogs” as a single category when you are trying to understand why some bite and others don't.

Cause or consequence?

A fair question remains: does *F. nucleatum* cause colorectal cancer, or does it merely exploit it?

The honest answer is that we do not yet know with certainty, and the truth may be that it does both. The strongest evidence suggests that *F. nucleatum* is a tumour accelerator rather than a tumour initiator. It does not appear to be sufficient, on its own, to transform a healthy colon cell into a cancerous one. What it does — through β -catenin activation, immune suppression, and pro-inflammatory signalling — is accelerate the growth and immune evasion of tumours that have already begun to form through the usual accumulation of driver mutations. Think of it not as the match that starts the fire but as the wind that fans it.

The clinical significance is not diminished by this distinction. *F. nucleatum*-enriched colorectal cancers are associated with worse prognosis, increased metastatic potential, and — in emerging data — resistance to certain chemotherapy regimens. The bacterium's presence in the tumour is a biomarker of aggressive disease, whether or not it is a cause. And the possibility that targeting

F. nucleatum — with antibiotics, with phage therapy, or with immune interventions that block its immunosuppressive effects — could slow tumour progression is being actively pursued in early-phase clinical trials.

28.3 The Microbiome as a Modulator of Chemotherapy

An accidental discovery

For decades, cancer pharmacology treated the gut microbiome as irrelevant. Chemotherapy drugs were designed to target rapidly dividing cells, and their efficacy was understood in terms of the drug's interaction with the tumour: DNA crosslinking, topoisomerase inhibition, antimetabolite competition. The billions of bacteria in the patient's gut were not part of the equation.

In November 2013, two studies published simultaneously in *Science* overturned this assumption so completely that the field has not looked back.

The first, by Norihiro Iida and colleagues, examined the platinum compound **oxaliplatin** — a mainstay of colorectal cancer treatment — in tumour-bearing mice [REF:iida2013]. They found that germ-free mice, or mice whose microbiomes had been depleted by broad-spectrum antibiotics, responded dramatically less well to oxaliplatin than conventionally colonised mice. The mechanism was revealing: oxaliplatin's anticancer effect depends, in part, on the generation of reactive oxygen species (ROS) by tumour-infiltrating myeloid cells. In the absence of a microbiome, these myeloid cells produced less ROS, and the drug's cytotoxic effect was blunted. The microbiome, it turned out, was providing the inflammatory “kindling” — the baseline innate immune activation, the tonic signal we discussed in Chapter 12 — that the drug needed to ignite its full anticancer effect.

The second study, by Sophie Viaud and colleagues in Laurence Zitvogel's group, told a complementary story about **cyclophosphamide**, one of the oldest and most widely used chemotherapy drugs [REF:viaud2013]. Cyclophosphamide damages the intestinal epithelium — an unfortunate but inevitable side effect — and this damage allows specific gram-positive bacteria, including *Lactobacillus johnsonii*, *Lactobacillus murinus*, and *Enterococcus hirae*, to translocate from the gut lumen into the mesenteric lymph nodes and the spleen. There, these bacteria stimulate the generation of pathogenic Th17 cells and memory Th1 responses — exactly the type of anti-tumour immune responses that cyclophosphamide is known to promote. In germ-free mice, or in mice treated with antibiotics that eliminated gram-positive bacteria, the drug's anticancer efficacy was significantly reduced.

The implication was startling. Two of the most important classes of chemotherapy — platinum compounds and alkylating agents — depend, for their full efficacy, on the presence and composition of the gut microbiome. The bacteria are not passive bystanders. They are active participants in the therapeutic mechanism.

The TIMER framework

Subsequent research, synthesised in a 2017 review by Alexander and colleagues in *Nature Reviews Gastroenterology & Hepatology*, established that the gut microbiome interacts with cancer drugs through at least five mechanisms, captured in the mnemonic **TIMER**: Translocation, Immunomodulation, Metabolism, Enzymatic degradation, and Reduced diversity [REF:alexander2017].

Translocation is the Viaud mechanism: bacteria crossing the damaged gut barrier and stimulating systemic immune responses that contribute to the drug's anticancer effect.

Immunomodulation is the lida mechanism: the microbiome maintaining the baseline activation state of innate immune cells that the drug requires to function.

Metabolism works in both directions. Gut bacteria can metabolically activate prodrugs — converting inactive compounds into their active forms — or they can inactivate drugs before they reach the tumour. The classic example is **irinotecan** (CPT-11), a topoisomerase I inhibitor used in colorectal cancer. Irinotecan is detoxified in the liver and excreted into the gut as an inactive glucuronide. But bacterial β -glucuronidases in the colon strip off the glucuronide, reactivating the drug in the gut lumen and causing severe diarrhoea — one of irinotecan's most feared side effects. Inhibiting bacterial β -glucuronidase with a targeted small molecule can prevent this reactivation without affecting the drug's anticancer activity, offering a microbiome-sparing approach to toxicity management [REF:wallace2010].

Enzymatic degradation is the flip side: bacteria that break down the drug before it can act. *Mycoplasma hyorhina* and certain *Gammaproteobacteria* carry the enzyme cytidine deaminase, which degrades **gemcitabine** — a nucleoside analogue used in pancreatic and lung cancer — into an inactive metabolite. Tumours that harbour these bacteria respond less well to gemcitabine, and this may be one reason why pancreatic cancers, which often contain intratumoral bacteria, show such variable responses to the drug [REF:geller2017].

Reduced diversity is the damage the drug inflicts on the microbiome itself. Chemotherapy, like antibiotics, depletes the gut community, and the resulting dysbiosis can create a vicious cycle: reduced diversity impairs immune function, which impairs the anti-tumour response, which requires dose escalation, which further depletes the microbiome.

The TIMER framework matters because it reframes the relationship between cancer pharmacology and the microbiome from “unrelated” to “deeply entangled.” Every chemotherapy drug passes through a patient who carries a microbial ecosystem, and that ecosystem modifies the drug's efficacy, its toxicity, and its side effects. The clinical implication — not yet realised in routine practice, but increasingly advocated — is that the microbiome should be considered a variable in drug dosing, treatment selection, and toxicity management, much as hepatic enzyme polymorphisms already are in the field of pharmacogenomics.

28.4 Immune Checkpoint Inhibitors: The Microbiome Predicts Response

A story already begun

We introduced this story in Chapter 12 — the 2015 Sivan mouse study, the 2018 trifecta of Gopalakrishnan, Matson, and Routy, the FMT trials of Baruch and Davar — and we will not retread that ground here. The reader who has come through Volume 1 already knows the core finding: the composition of the gut microbiome predicts whether a cancer patient will respond to immune checkpoint inhibitor (ICI) therapy, and transferring the microbiome from a responder to a non-responder can, in some cases, convert the non-responder into a responder.

What Chapter 12 could not do, because it was building the reader's understanding of the immune system itself, was place this finding in its broader oncological context. That is what this section provides.

The reproducibility challenge

The first and most persistent problem in the ICI-microbiome field has been inconsistency across cohorts. Gopalakrishnan's MD Anderson study pointed to *Faecalibacterium* as a marker of response. Matson's Chicago study pointed to *Bifidobacterium longum*. Routy's multinational study pointed to *Akkermansia muciniphila*. These are different organisms, living in different ecological niches, with different metabolic outputs. How could they all be “the” predictor?

A 2022 cross-cohort meta-analysis by Lee and colleagues in *Nature Medicine*, integrating 16S and metagenomic data from over 700 melanoma patients across multiple continents, made progress on this question [REF:lee2022]. They found that the associations were real but population-specific: the bacterial species most strongly associated with ICI response varied by geography, by diet, and by baseline microbiome composition. There was no single universal microbial biomarker. Instead, the data supported a model in which *functional* features of the microbiome — the capacity to produce certain metabolites, to maintain barrier integrity, to provide tonic immune stimulation — mattered more than the identity of the specific species performing those functions. Different species can perform the same ecological role in different populations, a concept ecologists call **functional redundancy**. The microbiome-ICI field had rediscovered, in cancer patients, a principle that ecologists had understood for decades.

The fungal layer

The cross-kingdom theme that we explored in Chapter 26 has also reached the ICI field. A 2025 multi-cohort meta-analysis of 976 faecal metagenomes across eight cohorts — spanning melanoma, non-small cell lung cancer, and renal cell carcinoma — examined not just bacteria but fungi and found that specific fungal species were independently associated with ICI efficacy [REF:fungal_ici2025]. The finding is preliminary, but it aligns with the broader argument of this book: the gut ecosystem is not a bacterial monoculture, and restricting our analysis to bacteria is like studying a forest by counting only the trees while ignoring the fungi, insects, and soil organisms that determine whether the forest thrives.

From observation to intervention

The ICI-microbiome field has moved rapidly from “interesting association” to “therapeutic target.” By 2026, a systematic review and meta-analysis identified thirty-six studies (twenty-five trials and cohorts, with a combined 2,746 patients) evaluating microbiome-modulating strategies in cancer patients receiving ICIs [REF:msit_meta2026]. The interventions ranged from FMT to probiotics to dietary modifications, and the pooled objective response rate (ORR) for microbiome modulation plus ICIs was forty per cent — a clinically meaningful figure in populations where many patients had previously failed standard immunotherapy.

The most promising approaches appear to be FMT from carefully selected donors and defined bacterial consortia — curated mixtures of specific species chosen for their immunostimulatory properties. Single-species probiotics have shown more modest effects, consistent with the ecological logic we developed in Chapter 23: adding one species to a complex ecosystem is a weak intervention compared with transplanting an entire community or a rationally designed subset of it.

Several pharma and biotech companies are now in late-phase clinical trials with defined consortia designed specifically to enhance ICI response. Whether these will reach routine clinical practice in the next five years is uncertain, but the trajectory is clear: the microbiome is no longer merely

a predictor of immunotherapy outcomes. It is becoming a modifiable variable in the treatment algorithm.

28.5 Predicting Cancer Recurrence: The Microbial Fingerprint

The problem of generalisability

In Chapter 12, we introduced the 2026 study by Usyk, Li, and colleagues — the Ahn group at NYU Langone — which analysed the gut microbiomes of 674 melanoma patients from the global Check-Mate 915 trial and found that specific bacterial taxa predicted cancer recurrence with 83 to 94 per cent accuracy, provided patients were first grouped by “microbial fingerprint” [REF:usyk2026]. We described the garden analogy, the concept of ecological matching, and the stability of the microbiome during immunotherapy. Here we place that study in the wider landscape of microbiome-based cancer prognostics.

Why fingerprinting changes everything

The Ahn study’s methodological innovation — matching patients by the overall similarity of their gut microbiomes before looking for disease-associated taxa — solved a problem that had haunted the field since the 2018 trifecta. Previous studies had tried to identify universal microbial biomarkers: species X predicts response, species Y predicts recurrence, regardless of who the patient is or where they live. These signatures worked well within the cohort in which they were discovered but failed, or performed inconsistently, when applied to patients from different geographic regions, different dietary backgrounds, or different ethnic groups.

The fingerprinting approach accepts the ecological reality that baseline microbiome composition varies. Instead of fighting this variation, it embraces it. By first sorting patients into ecologically similar groups — people whose guts look alike, regardless of their postal address — and then looking for predictive signatures within each group, the method produces predictions that are both accurate and portable. A signature discovered in North American patients transfers to European patients, but only to those European patients whose microbiome fingerprint matches the North American discovery cohort.

This is not merely a statistical trick. It reflects a genuine biological insight: the relationship between gut bacteria and cancer recurrence is mediated by the ecological context in which those bacteria operate. A species that is protective in one ecological setting may be neutral or even harmful in another, depending on the metabolic interactions, competitive relationships, and immune stimulation patterns that define that community. The fingerprint is the context.

Beyond melanoma

Melanoma has been the proving ground for microbiome-ICI research, largely because checkpoint inhibitors were first approved for melanoma, and the earliest cohorts with stool samples and clinical outcomes were melanoma cohorts. But the principles established in melanoma are now being extended to other cancers. Studies in non-small cell lung cancer, renal cell carcinoma, and hepatobiliary cancers have all reported microbiome associations with ICI response and survival, and the cross-cohort meta-analyses described in section 28.4 suggest that the microbiome’s predictive power may be a general feature of immune-mediated cancer therapy rather than a melanoma-specific finding.

The clinical scenario that the Ahn study anticipates is one in which a newly diagnosed cancer patient provides a stool sample alongside the usual blood draws and imaging at diagnosis. The microbiome is profiled — a process that takes days with current technology, not weeks — and the result informs treatment selection: patients whose microbiome predicts poor immunotherapy response might be offered alternative regimens or microbiome-modulating interventions before immunotherapy begins. Patients whose microbiome predicts high recurrence risk might receive more intensive surveillance or more aggressive adjuvant therapy. This is not speculative science fiction. The technology exists. What remains is the clinical validation — the large, randomised trials that demonstrate, prospectively, that microbiome-informed treatment decisions improve patient outcomes.

28.6 Could We Use the Microbiome for Early Cancer Detection?

The screening imperative

Cancer is far more treatable when caught early. This is not a platitude; it is a statistical fact with enormous clinical consequences. The five-year survival rate for colorectal cancer diagnosed at stage I is over ninety per cent; at stage IV, it is roughly fifteen per cent. For pancreatic cancer, the disparity is even starker: thirty to forty per cent at stage I, less than five per cent at stage IV. The value of a reliable early-detection screen — one that is non-invasive, inexpensive, and acceptable to the general population — is difficult to overstate.

Colonoscopy remains the gold standard for colorectal cancer screening, with a detection rate of approximately ninety-four per cent. But colonoscopy is expensive, requires bowel preparation, carries a small but real risk of complications, and — perhaps most importantly — is profoundly unpopular with patients. Compliance rates for colonoscopy screening programmes hover around sixty to sixty-five per cent in most countries. Many cancers that could be caught early are missed simply because the patient did not want to undergo the procedure.

The faecal immunochemical test (FIT), which detects blood in the stool, is a simpler alternative — but its sensitivity for early-stage cancers is limited, typically in the range of seventy to eighty per cent, and it misses a substantial proportion of pre-cancerous adenomas. The clinical need for something better — as accurate as colonoscopy, as simple as a stool sample — has driven an intense search for novel biomarkers. And the microbiome, it turns out, may be one of the best places to look.

The microbial signature of early cancer

The logic is straightforward. If a developing tumour alters the local microbial environment — through inflammation, through the release of tumour-derived nutrients, through changes in the mucosal surface — then those alterations should be detectable in the faecal microbiome before the tumour is large enough to produce symptoms or even to shed detectable amounts of blood. The microbiome, in effect, could serve as a sentinel — an early-warning system that responds to tissue changes that are too subtle for conventional diagnostics to detect.

Multiple studies have now demonstrated that the faecal microbiome of patients with colorectal cancer differs from that of healthy controls in ways that are statistically robust and, to some extent, reproducible across cohorts. The signature includes enrichment of species we have already discussed — *F. nucleatum*, *Bacteroides fragilis*, *Peptostreptococcus stomatis*, *Parvimonas micra* — and depletion of protective butyrate producers. Machine learning classifiers trained on these microbial

features can distinguish CRC patients from healthy controls with area-under-the-curve (AUC) values of 0.80 to 0.85 in cross-validation studies — respectable, but not yet at the threshold required for population-level screening.

The subspecies breakthrough

In 2025, a team at the University of Geneva led by Mirko Trajkovski made a methodological advance that may shift the field. Rather than analysing the microbiome at the species level — the resolution used by most previous studies — they constructed the first comprehensive catalogue of human gut bacteria at the **subspecies** level, identifying functional subgroups within species that had previously been treated as homogeneous [REF:trickovic2025].

The choice of resolution turned out to be critical. When they applied machine learning to subspecies-level microbial patterns in stool samples, their classifier detected ninety per cent of colorectal cancers — nearly matching colonoscopy's ninety-four per cent detection rate and substantially outperforming FIT and earlier microbiome-based classifiers [REF:trickovic2025]. The improvement came from the same insight that drove the Bullman clade study and the Ahn fingerprint study: the species level is too coarse. Within a single named species, there are functionally distinct subgroups — clades, subspecies, ecotypes — that behave differently in disease. Collapsing them into a single category obscures the signal. Splitting them apart reveals it.

A clinical trial is being set up at the Geneva University Hospitals to validate these findings prospectively and to determine which cancer stages and precancerous lesions the test can reliably detect. If the results hold up, the implications are profound: a simple stool sample, analysed by machine learning at subspecies resolution, could serve as a non-invasive screen with accuracy approaching that of colonoscopy, at a fraction of the cost and with none of the discomfort.

Beyond colorectal cancer

The microbiome-based screening concept is not limited to colorectal cancer. Emerging work has identified microbial signatures — both in stool and in circulating microbial DNA (cmDNA) in plasma — that can distinguish patients with pancreatic adenocarcinoma from healthy controls, with AUC values exceeding 0.90 in some studies. A 2024 study from the University of Chicago demonstrated that modification patterns in microbiome-derived cell-free RNA could predict colorectal cancer with ninety-five per cent accuracy — a liquid biopsy approach that combines the sensitivity of molecular diagnostics with the biological specificity of the microbiome [REF:cfrna_crc2024].

These are early results, from single-centre or small-cohort studies, and the road from promising biomarker to validated clinical screen is long. But the direction is clear. The microbiome is not only a modulator of cancer biology and a predictor of treatment response. It may also become one of the tools we use to find cancers before they kill.

A Relationship Still Unfolding

The microbiome's role in cancer is not a single story. It is at least four — carcinogenesis, treatment modulation, prognostication, and early detection — and each is at a different stage of maturity.

The *H. pylori* story is the most complete: from discovery to Nobel Prize to population-level screening programmes, spanning four decades of research and culminating in a genuine public health intervention that has saved millions of lives. The *F. nucleatum* story is a decade old and still working out its mechanistic details, but it has already changed how pathologists interpret tumour tissue and how oncologists think about the tumour microenvironment. The immunotherapy story — the realisation that the ecosystem in your gut determines whether a drug can kill a cancer in your skin — may be the most conceptually stunning finding in modern oncology, and the therapeutic implications are only beginning to be explored.

The early detection story is the youngest and the most uncertain. The technical barriers — reproducibility across populations, standardisation of methods, prospective validation — are real, and the history of cancer biomarker research is littered with promising candidates that failed in large-scale trials. But the logic is sound, the preliminary data are encouraging, and the clinical need is enormous.

What unites all four stories is a principle that this book has argued from its opening pages: the microbiome is not a passenger in human biology. It is a participant. It shapes the landscape in which disease arises, modifies the tools we use to fight disease, and offers information about disease that we are only beginning to learn how to read. Cancer, for all its genetic complexity, unfolds in an ecological context — and the ecology of the host includes the trillions of organisms that live within it.

In the next chapter, we turn to a different set of diseases — neurological and psychiatric — where the microbiome's role is more contested, the mechanisms are less clear, and the clinical applications are further from the bedside. But the underlying question is the same: how does the community we carry influence the diseases we develop? The answer, as we are learning, is: more than anyone suspected.

Chapter 28 — Key References

- [REF:hatakeyama2019] Hatakeyama, M. (2019). Malignant *Helicobacter pylori*-associated diseases: gastric cancer and MALT lymphoma. *Advances in Experimental Medicine and Biology*, 1149, 135–149. DOI: 10.1007/5584_2019_363
- [REF:raju2019] Raju, D., Hussey, S., Ang, M., Terebiznik, M.R., Bhatt, R., Bhatt, S., et al. (2012). Vacuolating cytotoxin and variants in *Helicobacter pylori*-associated VacA membrane channels. *Journal of Biological Chemistry*, 287(38), 31922–31934. DOI: 10.1074/jbc.M112.398990
- [REF:choi2018] Choi, I.J., Kim, C.G., Lee, J.Y., Kim, Y.-I., Kook, M.-C., Park, B., et al. (2018). Family history of gastric cancer and *Helicobacter pylori* treatment. *New England Journal of Medicine*, 382(5), 427–436. DOI: 10.1056/NEJMoa1909666
- [REF:castaño-rodríguez2017] Castaño-Rodríguez, N., Goh, K.-L., Fock, K.M., Mitchell, H.M. & Kaakoush, N.O. (2017). Dysbiosis of the microbiome in gastric carcinogenesis. *Scientific Reports*, 7, 15957. DOI: 10.1038/s41598-017-16002-5
- [REF:eradication_microbiome2025] Cheng, Y., Li, X., Zhang, W., et al. (2025). The impact of *Helicobacter pylori* infection and eradication therapies on gut microbiota: a systematic review of microbial dysbiosis and its implications in gastric carcinogenesis. *Frontiers in Cellular and Infection Microbiology*, 15, 1592977. DOI: 10.3389/fcimb.2025.1592977

- [REF:castellarin2012] Castellarin, M., Warren, R.L., Freeman, J.D., Dreolini, L., Krzywinski, M., Strauss, J., et al. (2012). *Fusobacterium nucleatum* infection is prevalent in human colorectal carcinoma. *Genome Research*, 22(2), 299–306. DOI: 10.1101/gr.126516.111
- [REF:kostic2012] Kostic, A.D., Gevers, D., Pedamallu, C.S., Michaud, M., Duke, F., Earl, A.M., et al. (2012). Genomic analysis identifies association of *Fusobacterium* with colorectal carcinoma. *Genome Research*, 22(2), 292–298. DOI: 10.1101/gr.126573.111
- [REF:abed2016] Abed, J., Emgård, J.E.M., Zamir, G., Faroja, M., Almogy, G., Grenov, A., et al. (2016). Fap2 mediates *Fusobacterium nucleatum* colorectal adenocarcinoma enrichment by binding to tumor-expressed Gal-GalNAc. *Cell Host & Microbe*, 20(2), 215–225. DOI: 10.1016/j.chom.2016.07.006
- [REF:fap2_structure2025] Tomida, S., Matsubara, T., Inoue, H., et al. (2025). Structural basis of *Fusobacterium nucleatum* adhesin Fap2 interaction with receptors on cancer and immune cells. *Nature Communications*, 16, 5012. DOI: 10.1038/s41467-025-63451-w
- [REF:rubinstein2013] Rubinstein, M.R., Wang, X., Liu, W., Hao, Y., Cai, G. & Han, Y.W. (2013). *Fusobacterium nucleatum* promotes colorectal carcinogenesis by modulating E-cadherin/ β -catenin signaling via its FadA adhesin. *Cell Host & Microbe*, 14(2), 195–206. DOI: 10.1016/j.chom.2013.07.012
- [REF:gur2015] Gur, C., Ibrahim, Y., Isaacson, B., Yamin, R., Abed, J., Gamliel, M., et al. (2015). Binding of the Fap2 protein of *Fusobacterium nucleatum* to human inhibitory receptor TIGIT protects tumors from immune cell attack. *Immunity*, 42(2), 344–355. DOI: 10.1016/j.immuni.2015.01.010
- [REF:bullman2024] Bullman, S., Pedamallu, C.S., Perna, F., et al. (2024). A distinct *Fusobacterium nucleatum* clade dominates the colorectal cancer niche. *Nature*, 628, 424–432. DOI: 10.1038/s41586-024-07182-w
- [REF:iida2013] Iida, N., Dzutsev, A., Stewart, C.A., Smith, L., Bouladoux, N., Weingarten, R.A., et al. (2013). Commensal bacteria control cancer response to therapy by modulating the tumor microenvironment. *Science*, 342(6161), 967–970. DOI: 10.1126/science.1240527
- [REF:viaud2013] Viaud, S., Saccheri, F., Mignot, G., Yamazaki, T., Daillère, R., Hannani, D., et al. (2013). The intestinal microbiota modulates the anticancer immune effects of cyclophosphamide. *Science*, 342(6161), 971–976. DOI: 10.1126/science.1240537
- [REF:alexander2017] Alexander, J.L., Wilson, I.D., Teare, J., Marchesi, J.R., Nicholson, J.K. & Kinross, J.M. (2017). Gut microbiota modulation of chemotherapy efficacy and toxicity. *Nature Reviews Gastroenterology & Hepatology*, 14(6), 356–365. DOI: 10.1038/nr-gastro.2017.20
- [REF:wallace2010] Wallace, B.D., Wang, H., Lane, K.T., Scott, J.E., Orans, J., Koo, J.S., et al. (2010). Alleviating cancer drug toxicity by inhibiting a bacterial enzyme. *Science*, 330(6005), 831–835. DOI: 10.1126/science.1191175
- [REF:geller2017] Geller, L.T., Barzily-Rokni, M., Danino, T., Jonas, O.H., Shental, N., Nejman, D., et al. (2017). Potential role of intratumor bacteria in mediating tumor resistance to the chemotherapeutic drug gemcitabine. *Science*, 357(6356), 1156–1160. DOI: 10.1126/science.aah5043

- [REF:lee2022] Lee, K.A., Thomas, A.M., Bolte, L.A., Björk, J.R., de Ruijter, L.K., Armanini, F., et al. (2022). Cross-cohort gut microbiome associations with immune checkpoint inhibitor response in advanced melanoma. *Nature Medicine*, 28(3), 535–544. DOI: 10.1038/s41591-022-01695-5
- [REF:fungal_ici2025] Papadimitriou, K., Thomas, A.M., et al. (2025). Intestinal fungal signatures and their impact on immune checkpoint inhibitor efficacy: a multi-cohort meta-analysis. *npj Biofilms and Microbiomes*, 11, 827. DOI: 10.1038/s41522-025-00827-2
- [REF:msit_meta2026] Zheng, H., Li, M., et al. (2026). Impact of microbiome-modulating strategies in cancer patients receiving immunotherapy (MSIT): a systematic review and meta-analysis. *Scientific Reports*, 16, 44743. DOI: 10.1038/s41598-026-44743-7
- [REF:usyk2026] Usyk, M., Li, H., Osman, I., Weber, J.S., Knight, R., González, A., Hayes, R.B. & Ahn, J. (2026). Gut microbiome is associated with recurrence-free survival in patients with resected high-risk melanoma receiving adjuvant immune checkpoint blockade. *Cell*. DOI: 10.1016/j.cell.2026.03.041
- [REF:trickovic2025] Tričković, M., Kieser, S., Zdobnov, E.M. & Trajkovski, M. (2025). Subspecies-level resolution of the human gut microbiome for colorectal cancer detection. *Cell Host & Microbe*, 33(8). DOI: 10.1016/j.chom.2025.07.015
- [REF:cfrna_crc2024] University of Chicago research group (2024). Microbial RNA modifications predict colorectal cancer in liquid biopsy. *[Journal to be confirmed]*.
- [REF:sivan2015] Sivan, A., Corrales, L., Hubert, N., et al. (2015). Commensal *Bifidobacterium* promotes antitumor immunity and facilitates anti-PD-L1 efficacy. *Science*, 350(6264), 1084–1089. DOI: 10.1126/science.aac4255
- [REF:gopalakrishnan2018] Gopalakrishnan, V., Spencer, C.N., Nezi, L., et al. (2018). Gut microbiome modulates response to anti-PD-1 immunotherapy in melanoma patients. *Science*, 359(6371), 97–103. DOI: 10.1126/science.aan4236
- [REF:matson2018] Matson, V., Fessler, J., Bao, R., et al. (2018). The commensal microbiome is associated with anti-PD-1 efficacy in metastatic melanoma patients. *Science*, 359(6371), 104–108. DOI: 10.1126/science.aao3290
- [REF:routy2018] Routy, B., Le Chatelier, E., Derosa, L., et al. (2018). Gut microbiome influences efficacy of PD-1-based immunotherapy against epithelial tumors. *Science*, 359(6371), 91–97. DOI: 10.1126/science.aan3706
- [REF:baruch2021] Baruch, E.N., Youngster, I., Ben-Betzalel, G., Ortenberg, R., Lahat, A., Katz, L., et al. (2021). Fecal microbiota transplant promotes response in immunotherapy-refractory melanoma patients. *Science*, 371(6529), 602–609. DOI: 10.1126/science.abb5920
- [REF:davar2021] Davar, D., Dzutsev, A.K., McCulloch, J.A., Rodrigues, R.R., Chauvin, J.-M., Morrison, R.M., et al. (2021). Fecal microbiota transplant overcomes resistance to anti-PD-1 therapy in melanoma patients. *Science*, 371(6529), 595–602. DOI: 10.1126/science.abf3363

Chapter 29: Neurological and Psychiatric Disease

29.0 The Hardest Organ

The chapter you have just finished made a case that is, by the standards of microbiome science, unusually solid. In cancer, we can now point to specific bacteria that directly cause particular tumours, to microbial communities that decide whether an immunotherapy will work, and to stool-based tests that are beginning to look like genuine diagnostic tools. The evidence in oncology has moved from association to mechanism to clinical intervention over a single working lifetime.

The brain is different. The brain is, in fact, the hardest organ for the microbiome story to reach. There are structural reasons for this. The blood–brain barrier — a near-continuous sheet of tightly joined endothelial cells wrapped in astrocytic feet — attenuates most of the direct microbial signals that elsewhere in the body would flood tissues at will. The brain has no lymphatic system in the classical sense, no resident bacterial community of any substantial size, and a cellular population that turns over, for practical purposes, almost not at all after childhood. It is slow to change, hard to biopsy in life, and — most importantly for a translation effort — hard to model in laboratory animals. The mouse brain has many things the human brain has; it also lacks a great deal, including a prefrontal cortex of comparable elaboration, the capacity for declarative language, and the cultural scaffolding that shapes much of human cognition and mood.

So when a researcher claims that gut microbes are driving depression, or autism, or Parkinson's disease, or Alzheimer's, they are making a claim that has to climb up a very tall ladder before it lands in the clinic. A signalling molecule produced in the colon has to either cross the blood–brain barrier in sufficient quantity, or signal through a nerve (usually the vagus), or change a peripheral physiology (inflammation, for example) in a way that eventually reaches the brain. The end-point of that journey — a depressive episode, a tremor, a memory lapse — then has to be measurable reliably enough that a clinical trial can show an effect. Each step of this chain loses signal. By the time you reach the mouse model of human depression or the first-in-human probiotic trial, what survives is often a whisper.

This chapter takes four whispers seriously. The first is the loudest: the story of Parkinson's disease, where the Braak hypothesis that we introduced in Chapter 11 has matured over the past decade into one of the most compelling examples of microbial biology reaching from gut to brain. The second is Alzheimer's disease, where the infection hypothesis of neurodegeneration — also introduced in Chapter 11 — has just survived its first large clinical trial, and where the pattern of strong observational evidence meeting disappointing randomised results is being rerun with microbial hypotheses. The third is depression, where the gut–brain axis framework we built in Chapter 13 has begun to yield small but real clinical effects from microbiome-targeted interventions, and where the attempt to extend from mouse to patient is now sufficiently advanced that we can see what has worked and what has not. Woven in are shorter treatments of autism spectrum disorder, which sits at the edge of the clinical evidence frontier, and of multiple sclerosis and amyotrophic lateral sclerosis, where microbial contributions are now being actively dissected but where the picture is still genuinely early.

The chapter's fourth and longest thread is not a disease but a method. Section 29.4 is a deliberately critical pass through the translation problem itself — a working explanation of why so many exciting mouse findings fail in humans, what germ-free mice are and are not, what mouse behavioural tests actually measure, and what kind of trial would be needed before any microbiome-targeted intervention could responsibly enter routine neuropsychiatric practice. None exists yet. The purpose of that section is not to deflate the reader's enthusiasm — the science here is genuinely promising — but to equip them to read the next wave of headlines with the calibrated scepticism that good science demands of itself.

One other note on tone. In the cancer chapter we could, with some confidence, describe interventions that are already changing clinical practice. In this chapter we cannot. Almost everything described here lives in the zone between “demonstrated in animals” and “associated in humans.” The handful of randomised trials that exist are small, single-centre, or have delivered only modest improvements. That is not a reason to dismiss the field. It is a reason to read it carefully.

29.1 Parkinson's Disease: Does It Start in the Gut?

Parkinson's disease is, at first glance, an unlikely candidate for a gut-based origin story. Its cardinal features — resting tremor, muscular rigidity, slowness of movement, postural instability — are movement symptoms that arise from the loss of dopamine-producing neurons in a small midbrain structure called the substantia nigra. The microscopic hallmark of the disease is the Lewy body, an intracellular aggregate whose principal component is a misfolded form of the protein α -synuclein. The most effective treatment we have, levodopa, replaces the dopamine those dying neurons are no longer making. By any reasonable reading, Parkinson's disease is a disease of the brain.

And yet. Patients with Parkinson's disease, when asked carefully, will tell you that their symptoms did not begin with the tremor. They will tell you that they had constipation for years — sometimes decades — before the movement problems appeared. They will tell you that their sense of smell faded. They will tell you that their partner noticed them kicking and shouting in their sleep, acting out dreams in a way that is now recognised as a specific disorder: REM sleep behaviour disorder, or RBD. Longitudinal follow-up of people with idiopathic RBD shows that more than 80 per cent of them will go on to develop Parkinson's disease, or the related synucleinopathies dementia with Lewy bodies or multiple system atrophy, within 15 years of the RBD diagnosis [REF:postuma2019]. Constipation precedes motor Parkinson's in around 60 per cent of cases, often by a decade or more. The prodromal phase of Parkinson's — the period between earliest disease and first motor symptoms — is now thought to span 10 to 20 years.

If the disease begins that long before the tremor, and if its earliest manifestations are autonomic and sensory rather than motor, it is reasonable to ask where it actually starts. In Chapter 11 we met the German neuroanatomist Heiko Braak and his hypothesis that Parkinson's disease begins, in many cases, not in the substantia nigra but in two peripheral locations: the olfactory bulb and the enteric nervous system — the dense mesh of neurons embedded in the wall of the gut, which we sometimes call “the second brain.” From those beginnings, misfolded α -synuclein propagates upward along the vagus nerve, reaching the dorsal motor nucleus of the vagus in the brainstem and then climbing through the midbrain toward the substantia nigra. By the time the characteristic motor symptoms appear, the pathology has been moving through the body for years.

This hypothesis was, when Braak first proposed it, controversial. It is now mainstream. A refine-

ment — the body-first versus brain-first framework proposed by Per Borghammer and colleagues — divides Parkinson's into two variants with rather different clinical signatures [REF:horsager2020]. In body-first Parkinson's, α -synuclein pathology begins in the enteric or autonomic nervous system and ascends to the brainstem before reaching the substantia nigra. These patients tend to have prominent early autonomic symptoms, constipation, orthostatic hypotension, and — strikingly — RBD. In brain-first Parkinson's, the pathology appears to begin somewhere in the central nervous system and spread downward. These patients have fewer gastrointestinal symptoms and less RBD. The microbiome is most plausibly relevant to the body-first variant, which may account for something between a third and a half of cases. If that sub-classification proves robust, we may find that the gut–brain story fits a specific subset of Parkinson's patients beautifully while saying little about the rest — a pattern the reader will by now recognise as typical of microbiome medicine.

The two pillars

Two papers, a year apart and from opposite sides of the Atlantic, form the foundation of the modern Parkinson's–microbiome literature. Together they established that there is a real signal, that it is at least partly causal in mice, and that it deserves the decade of follow-up work that has since been poured into it.

The first came from Filip Scheperjans and colleagues at the University of Helsinki in 2015 [REF:scheperjans2015]. They sequenced the 16S rRNA gene from stool samples of 72 patients with Parkinson's disease and 72 matched controls — a modest cohort by the standards of the field today, but the first carefully controlled human study of its kind. The headline finding was dramatic: the Prevotellaceae, a family of gut bacteria that ferments dietary fibre to short-chain fatty acids, was reduced by 77.6 per cent in Parkinson's patients compared with controls. A second finding would prove prognostic: the abundance of the Enterobacteriaceae, a family that includes *Escherichia coli* and several opportunistic pathogens, correlated with the severity of postural instability and gait difficulty. The signal was not subtle. Scheperjans' paper appeared in *Movement Disorders* and immediately opened the field.

The second came from Timothy Sampson, working in the laboratory of Sarkis Mazmanian at Caltech [REF:sampson2016]. Its mouse genetics were elegant. Sampson used a transgenic mouse strain that over-expresses human α -synuclein throughout the nervous system — an animal model that develops progressive motor deficits reminiscent of Parkinson's disease. When Sampson raised these mice as germ-free animals — in sterile isolators with no microbiota of any kind — they had nearly normal motor function and minimal microglial activation in the brain. When he colonised them with a standard laboratory mouse microbiota, the motor deficits returned. Feeding germ-free α -synuclein mice short-chain fatty acids — the fermentation products that gut bacteria normally supply — was sufficient to restore the pathology. And in the pivotal experiment, mice colonised with stool from Parkinson's patients developed worse motor deficits than mice colonised with stool from matched controls.

This was the experiment that moved the field from association to mechanism — or rather, moved it to mechanism in mice, which is not quite the same thing. Scheperjans had shown that the Parkinson's microbiome is different from the healthy one. Sampson had shown that, at least in a susceptible mouse, the microbiome is actively driving motor dysfunction. What remained — and remains — unresolved is whether the same mechanism operates in humans. Mouse models of α -synuclein toxicity are genetically stronger than the human disease and are probably not ideal surrogates for the decades-long human pathology. But Sampson's paper made it impossible to dismiss the micro-

biome as an epiphenomenon. Something in the community of bacteria living in the Parkinson's gut is actively relevant to the brain above it.

The metagenomic era

The decade since Scheperjans' paper has produced dozens of follow-up studies in cohorts around the world. Individually they have produced inconsistent results: the Prevotella finding has replicated in some cohorts and not in others; the specific taxa most strongly associated with disease differ between Finland and Shanghai and Alabama. But at the level of functional categories — the kinds of bacteria Parkinson's patients have more or less of — a broad pattern has emerged. Parkinson's microbiomes tend to have fewer producers of short-chain fatty acids (the butyrate-generating Firmicutes, including *Faecalibacterium prausnitzii* and *Roseburia* species), more mucin-degraders (*Akkermansia muciniphila*, curiously in the “too much” rather than “too little” direction that it usually occupies in other diseases), and more opportunistic pathogens (Enterobacteriaceae). The picture is one of a broadly disturbed ecosystem rather than a single pathogenic species.

The best illustration of this shift came in 2022, when Zachary Wallen and colleagues at the University of Alabama at Birmingham published the largest metagenomic study of Parkinson's disease to date [REF:wallen2022]. With 490 patients and 234 controls and deep shotgun sequencing — which reads fragments of all the genes in the community, not just the 16S ribosomal marker that earlier studies used — they had the statistical power to go beyond “this taxon is different” and ask, “which biological functions are different?” The answer was striking: more than 30 per cent of the species, genes, and metabolic pathways tested had altered abundances in Parkinson's. The Parkinson's gut was not subtly off; it was comprehensively remodelled. Wallen's team also showed that the disturbances clustered in networks — groups of taxa that rose and fell together — which is consistent with a disturbed ecosystem shifting to a new quasi-stable state rather than responding to a single provoking factor.

A 2024 Canadian study by Avril Metcalfe-Roach and colleagues extended this approach at species-level resolution, with 176 Parkinson's patients and 100 controls [REF:metcalferoach2024]. It reported species-specific and function-specific differences that correlated both with disease-related metabolites and with the rate at which patients' symptoms progressed. The idea of a Parkinson's microbiome not just as a marker of present disease but as a predictor of future trajectory is beginning to take shape.

Most intriguingly, the dysbiosis can be seen before the disease. Studies of patients with idiopathic RBD — those most likely to progress to Parkinson's — show microbiome signatures that resemble early-Parkinson's signatures more than they resemble healthy ones [REF:heintzbuschart2018]. Butyrate-producers are depleted; opportunists are enriched. First-degree relatives of Parkinson's patients, when they also show early prodromal features, carry similar signatures. This matters enormously for the causality question. If the microbiome changes only after the dopaminergic neurons start dying, it is an interesting consequence of disease. If it changes years before the first motor symptom, it has a chance to be part of the cause.

A note of caution is in order here. In 2025, a large machine-learning meta-analysis by a European consortium pooled microbiome data from nearly 4,500 Parkinson's patients across multiple cohorts [REF:romano2025]. Classifiers trained to distinguish Parkinson's from healthy controls within a single study achieved mean areas under the curve of around 72 per cent — decent but not outstanding. When the same classifiers were tested across cohorts, performance fell to about 61 per cent — barely better than a coin toss. The same generalisability problem we met in the cancer

microbiome story in Chapter 28 is present here. There is a real signal, but it is entangled with local diet, genetics, and ethnicity in ways that make a universal “Parkinson’s microbiome signature” more elusive than early papers suggested. Whether the microbial-fingerprint approach that Jiyoung Ahn’s group introduced for melanoma (Chapter 28) can rescue portability in Parkinson’s as well is an open question.

The appendix, α -synuclein, and the virtue of replication

In 2018, Bryan Killinger and colleagues at Van Andel Research Institute published a paper in *Science Translational Medicine* that electrified the field [REF:killinger2018]. Using two large epidemiological datasets covering more than 1.6 million individuals, they reported that surgical removal of the appendix — appendectomy — was associated with a 19–25 per cent lower risk of subsequently developing Parkinson’s disease, with the strongest effect in rural dwellers. They then examined the appendix itself and found that the healthy appendix is loaded with α -synuclein, including truncated forms of the protein that have the characteristic of being able to seed the fibrillar aggregates that form Lewy bodies. Appendix lysates, applied to α -synuclein monomers in a test tube, could nucleate their misfolding.

The interpretation was seductive. The appendix, a previously vestigial-seeming organ, might be one of the original sites at which α -synuclein pathology initiates in the body-first variant of Parkinson’s. Remove the appendix, and you lower the risk of the disease. Killinger’s paper was widely covered in the press and widely cited in subsequent reviews.

And then, awkwardly, the effect started to fade. Subsequent epidemiological studies in Swedish, Korean, and Taiwanese registers produced mixed or null results. A 2025 meta-analysis in *Frontiers in Neurology*, pooling the evidence across cohorts, concluded that there was no robust association between appendectomy and Parkinson’s risk [REF:appendectomy_meta2025]. The α -synuclein finding — that the appendix contains seeding-competent forms of the protein — has replicated. The epidemiological protective effect of its removal has not.

How do we reconcile this? The honest answer is that a striking association in a single large dataset often shrinks or disappears when tested elsewhere. Epidemiological associations are prone to confounding by health-care access (who gets an appendectomy? people who see doctors), by smoking (inversely associated with Parkinson’s through mechanisms still debated), and by the ethnic and geographic heterogeneity of the populations surveyed. Killinger’s paper is not wrong to have found a signal in his data. It is, however, wrong to have claimed a causal protective effect of appendectomy on the basis of that data alone. The appendix contains α -synuclein seeds; whether removing it protects against Parkinson’s is, at the time of writing, unresolved.

This is a useful teaching moment, and the reader who has followed the book to this point will recognise the pattern. A striking first paper, a great press cycle, a decade of mixed follow-up, and a final meta-analytic verdict of “maybe nothing.” This is how medical science often works. It is not a scandal; it is the correction mechanism doing its job. But it should temper our confidence in every new microbiome-and-brain headline we read.

The FMT trials

If the Parkinson’s microbiome is genuinely disturbed, and if the disturbance is plausibly part of disease biology, then the most direct intervention is to replace the disturbed community with a healthy one. Faecal microbiota transplantation — whose principles we explored in detail in Chapter

24 — is the most powerful ecological tool we currently have. Two randomised controlled trials have asked whether it works in Parkinson’s disease.

The first, from a Chinese team led by Le-Jian Xue in 2023, was a single-centre randomised placebo-controlled trial of oral FMT in 54 patients with Parkinson’s [REF:xue2023]. At the end of the trial period, FMT-treated patients had improved autonomic symptoms, improved gastrointestinal symptoms, and greater microbial complexity than the placebo group. Motor improvements were modest. The effect sizes were real but small.

The second, and the one that moved the field, was the GUT-PARFECT trial from Arnout Bruggeman and colleagues at Ghent University in Belgium in 2024 [REF:bruggeman2024]. GUT-PARFECT was a phase 2, double-blind, placebo-controlled trial of a single nasoduodenal FMT in 46 patients with mild-to-moderate Parkinson’s disease. Patients were randomised to receive either donor stool or their own stool (re-instilled as placebo). They were assessed on the Movement Disorder Society Unified Parkinson’s Disease Rating Scale — the MDS-UPDRS — a standardised instrument that scores motor features from tremor to bradykinesia. The trial’s primary endpoint was the change in the motor score from baseline to 12 months.

At 12 months, the FMT group had improved, in the off-medication state, by 5.8 points on the motor scale. The placebo group had improved by 2.7 points. The between-group difference — a 3.1-point improvement favouring FMT — was statistically significant ($p = 0.0235$). Adverse events were limited to transient abdominal discomfort in the first few weeks after the procedure. A single FMT had produced a small but real and durable improvement in motor symptoms of Parkinson’s disease.

This is a striking result, and it deserves to be taken seriously. It is also a result that should be held with some caution. GUT-PARFECT was a single-centre trial with a small number of participants. A three-point improvement on the motor scale is real but clinically modest — it is roughly the difference between “a bad morning” and “a good morning” for a patient with Parkinson’s disease, and it is substantially less than the improvement one would see from adjusting a levodopa dose. The placebo condition — the patient’s own stool instilled via a nasoduodenal tube — was not plausibly blinded to the medical team, though efforts were made to blind the assessors. And the finding that a single FMT could produce a durable 12-month improvement is, frankly, surprising. The microbial community that FMT establishes typically drifts over weeks to months toward something resembling the recipient’s pre-treatment state, as we saw in Chapter 24. How a single dose produces a year-long motor benefit — if, indeed, it does — will need mechanistic dissection and, more urgently, independent replication in a multi-centre trial.

At the time of writing, several follow-up trials are in progress. Until they report, GUT-PARFECT is best understood as a promising proof of concept rather than a practice-changing result. It is, however, the first randomised evidence that a microbiome-targeted intervention can measurably modify motor outcomes in Parkinson’s disease. That alone is worth the attention the field has given it.

The levodopa twist

The microbiome does not only influence whether Parkinson’s develops. It also influences how well our best treatment for it — levodopa — works. In Chapter 22 we met Vayu Maini Rekda’s 2019 *Science* paper showing that certain strains of *Enterococcus faecalis* carry a tyrosine decarboxylase enzyme that can decarboxylate levodopa to dopamine in the gut, before the drug has a chance to reach the brain [REF:mainirekda2019]. Dopamine cannot cross the blood–brain barrier; levodopa

can. When a patient's gut bacteria strip the carboxyl group from levodopa in the lumen, the resulting peripheral dopamine is wasted, and the brain receives less of the drug than was intended. The clinical consequence is patient-to-patient variability in levodopa response that has puzzled neurologists for decades — variability that, we now know, has a microbial component.

The ecological logic here is almost comical. We prescribe levodopa to compensate for the loss of dopamine-producing neurons in the substantia nigra; and then a community of gut bacteria — whose composition is itself partly a consequence of the Parkinson's disease state — short-circuits our carefully designed molecule before it ever reaches its target. The strongest available inhibitor of this bacterial enzyme is an analogue of tyrosine itself, and efforts to develop it as a clinical adjunct to levodopa are under way. The broader message, introduced in Chapter 22 and revisited here, is that the microbiome sits between the patient and almost every drug we give them, modifying both the disease and its treatment in ways that our pharmacology has, until recently, mostly ignored.

The state of the Parkinson's story

Of the four disease stories in this chapter, Parkinson's is the strongest. The Braak hypothesis has matured. The prodromal phase has been characterised. Gut microbiome signatures have been replicated at the functional level across cohorts, despite unhelpful inconsistencies at the species level. A mouse model has provided mechanistic support — α -synuclein pathology is dependent on microbiota in at least one transgenic mouse. Two randomised controlled trials, one of them double-blinded, have shown modest improvements from FMT. Our best drug is modified by our gut community. The field has moved, over a single decade, from “is there anything here?” to “how do we turn this into a therapy?”

It is also the story that most clearly shows the limits of where the field currently stands. GUT-PARFECT's three-point motor improvement is real but not revolutionary. Generalisability of the microbiome signature across populations is poor. The causal arrow — does dysbiosis trigger α -synuclein pathology, or does α -synuclein pathology, by altering gut motility and autonomic function, produce dysbiosis? — is not settled even in mice. Progress in the next five years will depend on trials that are larger, multi-centre, and powered to stratify by body-first versus brain-first disease. What the field has demonstrated is that the question is worth asking. What remains to be demonstrated is that asking it will change patients' lives.

29.2 Alzheimer's Disease and Neuroinflammation

If Parkinson's disease is the microbiome's strongest neurological case, Alzheimer's disease is its most disorienting. The Alzheimer's literature has long been dominated by the amyloid cascade hypothesis — the idea that the misfolded peptide amyloid- β is the proximate driver of the disease, with its downstream consequences including the neurofibrillary tangles of tau, widespread synaptic failure, and eventual dementia. Three decades of drug development built around this hypothesis produced, for a long time, nothing but failures. The recent arrival of anti-amyloid antibodies — lecanemab and donanemab — has given the cascade hypothesis its first clinical wins, but the effect sizes remain modest and the drugs carry non-trivial risks of brain swelling and bleeding. Even the most optimistic readings of the current Alzheimer's landscape suggest that we are modifying the disease, not curing it.

Into this long-running controversy, the microbiome story arrives with a set of claims that are, in

some ways, orthogonal to the cascade debate and, in others, intimately intertwined with it. The claims fall under two headings. The first — developed in Chapter 11 as the infection hypothesis — proposes that chronic or latent infections by specific organisms contribute to Alzheimer's pathogenesis. The second, which Chapter 11 touched on only briefly, proposes that the composition of the gut microbiome shapes systemic inflammation and blood–brain barrier integrity in ways that accelerate or decelerate disease. This chapter extends both threads into their clinical-trial era.

The gut microbiome of Alzheimer's disease

The first careful characterisation of the Alzheimer's gut microbiome was published in 2017 by Nicholas Vogt and colleagues at the University of Wisconsin [REF:vogt2017]. The study — 25 Alzheimer's patients, 25 age- and sex-matched controls drawn from the Wisconsin Alzheimer's Disease Research Center — was small, but it was carefully done, and it established a pattern that subsequent larger studies have broadly reproduced. Compared with controls, Alzheimer's patients had reduced microbial richness and reduced diversity. At higher taxonomic levels, the phylum Firmicutes was reduced and Bacteroidetes was increased. At genus level, *Bifidobacterium* was depleted. Cerebrospinal-fluid biomarkers of Alzheimer's pathology — the ratio of A β 42 to A β 40, and phosphorylated tau — correlated with microbial abundances in several taxa.

A decade of follow-up work has produced a picture consistent with Vogt's. A 2024 systematic review and meta-analysis by Ana Jimenez-García and colleagues pooled the evidence across more than a dozen cohorts [REF:jimenezgarcia2024]. The consistent signal: depletion of short-chain fatty-acid producers — particularly *Faecalibacterium prausnitzii* and *Bifidobacterium* — and enrichment of Proteobacteria, including opportunistic genera such as *Escherichia* and *Shigella*. The signature overlaps substantially with that of Parkinson's disease, with depressive disorders, with inflammatory bowel disease, and with general “ill health.” It is not Alzheimer's-specific. This is an important point that we will return to in Section 29.4: the microbial changes that accompany Alzheimer's are largely non-specific indicators of disturbed gut ecology, and any specific aetiological role is hidden within a pattern that also accompanies many other conditions.

A more tantalising line of evidence comes from studies of preclinical Alzheimer's — individuals who are cognitively normal but carry positron emission tomography or cerebrospinal-fluid evidence of early amyloid or tau pathology. In 2023, a group at Washington University in St. Louis reported gut microbiome differences between preclinical-Alzheimer's and control participants that were detectable in people whose cognition was still intact [REF:ferreiro2023]. If this holds up — and the cohorts will need to be expanded — it suggests that the microbial signature appears early in the disease process, potentially years before the first symptom. This is the same prodromal-signal story we met in Parkinson's, and it carries the same causal potential. A signature that is present in the preclinical phase might be a candidate for early intervention. A signature that appears only in established disease is likely to be a consequence, not a cause.

The LPS–inflammation–barrier axis

How might the gut microbiome actually influence a disease whose principal pathology sits behind the blood–brain barrier? The dominant mechanistic model, assembled from findings that will by now be familiar, goes roughly as follows. Gut dysbiosis — however it arises — tends to thin the protective mucus layer and weaken the tight junctions of the intestinal epithelium. Bacterial components, particularly the lipopolysaccharide (LPS) of Gram-negative cell walls, leak into the portal circulation at low levels. This chronic, low-grade endotoxaemia drives systemic inflammation — raised

CRP, raised IL-6, raised TNF- α — which in turn degrades the tight junctions of the blood–brain barrier itself. Microbial products, pro-inflammatory cytokines, and peripheral immune cells cross into the brain parenchyma more readily. There they activate microglia, the brain’s resident immune cells, pushing them from their surveillance state toward a persistently activated, pro-inflammatory phenotype. Chronic microglial activation impairs the clearance of amyloid- β and may accelerate its aggregation; it also drives further tau pathology and direct synaptic damage.

This framework draws on mechanisms we covered in earlier chapters — the leaky gut of Chapter 12, the blood–brain barrier permeability demonstrated by Braniste and Pettersson in 2014 (Chapter 11), the antimicrobial role of A β itself (Eimer, Tanzi and Moir, Chapter 11). The 2023 review by Arabi and colleagues, introduced in Chapter 11, remains a good synthesis: it argues that the infection hypothesis and the amyloid cascade hypothesis of Alzheimer’s are not competing — they are two ends of a single pathway in which diverse microbial triggers converge on a shared neuroinflammatory response [REF:arabi2023]. If this framework is correct, the gut microbiome is one of several possible upstream contributors to chronic neuroinflammation in Alzheimer’s, alongside latent viral infections, chronic periodontitis, and — in some models — the intrinsic antimicrobial response to misfolded protein itself.

What the framework is not is a proof. LPS-driven inflammation is well established in peripheral tissues; the specific contribution of gut-derived LPS to human brain pathology is still being characterised. Short-chain fatty acids do tighten the blood–brain barrier in rodent models; the quantitative impact of gut-produced butyrate on a human adult’s brain barrier is unknown. The story is plausible, internally consistent, and partly supported by animal data. It is not yet a validated pathway for which the therapeutic implications are clear.

The atuzaginstat story: a cautionary tale

In Chapter 11 we met the 2019 paper from Stephen Dominy and the company Cortexyme reporting that the oral bacterium *Porphyromonas gingivalis* — the principal pathogen of chronic periodontitis — and its toxic proteases, the gingipains, could be detected in the brains of Alzheimer’s patients at autopsy [REF:dominy2019]. Dominy argued that gingipains were a driver of Alzheimer’s neurodegeneration, and Cortexyme developed a gingipain inhibitor, atuzaginstat (COR388), which entered a large phase 2/3 clinical trial — the GAIN trial — in patients with mild-to-moderate Alzheimer’s disease.

GAIN was the first registrational-scale test of an antimicrobial approach to Alzheimer’s. It was large: 643 participants across North America and Europe, randomised to atuzaginstat at one of two doses or to placebo, with cognitive and functional endpoints measured over 48 weeks. The result, reported in 2021, was a failure [REF:gain2021]. Atuzaginstat did not meet the primary cognitive endpoint (the Alzheimer’s Disease Assessment Scale – Cognitive Subscale 11, or ADAS-Cog 11) or the primary functional endpoint (the Alzheimer’s Disease Cooperative Study – Activities of Daily Living). On the pooled population, patients on atuzaginstat declined at about the same rate as patients on placebo.

But the trial did produce a subgroup signal. In the pre-specified subset of participants with detectable *P. gingivalis* DNA in saliva at baseline — about 37 per cent of the study population, 242 people — atuzaginstat-treated patients showed dose-dependent slowing of cognitive decline, with the higher dose associated with a 57 per cent reduction in the rate of decline relative to placebo. This is exactly what one would predict if gingipains were driving disease specifically in patients with active *P. gingivalis* infection.

It is also exactly the kind of post-hoc rescue that a failed clinical trial often produces, and which, in the drug development literature, has a poor track record of replicating in confirmatory studies. The Food and Drug Administration placed a partial clinical hold on the atuzaginstat programme because of liver safety concerns. Cortexyme restructured, rebranded, and eventually abandoned the programme. Whether a gingipain inhibitor — or another *P. gingivalis*-targeted agent — could succeed in a better-designed trial, enrolled prospectively in *P. gingivalis*-positive participants, remains an open question. But the simple claim that removing a specific oral bacterium would slow Alzheimer's disease has, at present, failed its first major clinical test.

This is the pattern to watch for in the next decade. Microbiome-inspired trials in Alzheimer's disease will increasingly enter the clinic; they will, at least at first, tend to fail their primary endpoints; and they will often reveal suggestive subgroup signals that may or may not survive confirmation. The AD field has been through this cycle many times with amyloid-targeted drugs, and it has learned, painfully, that post-hoc subgroup rescues are usually how a programme dies gracefully rather than how it succeeds. Microbial hypotheses will have to run the same gauntlet.

Antivirals, HSV-1, and a large observational signal

A very different kind of evidence comes from population data on herpes simplex virus and dementia. In Chapter 11 we met the long-running hypothesis that latent HSV-1 infection, reactivating periodically in neurons of the temporal lobe, contributes to the amyloid and tau pathology of Alzheimer's disease. That hypothesis received striking epidemiological support in 2018 from a nationwide Taiwanese cohort study by Nian-Sheng Tzeng and colleagues [REF:tzeng2018]. Using the Taiwan National Health Insurance Research Database, Tzeng's team identified 8,362 individuals with newly diagnosed HSV infection and 25,086 matched non-infected controls, and followed them for a decade. HSV infection was associated with a 2.56-fold increased risk of dementia (adjusted hazard ratio). Among those with HSV, anti-herpetic treatment — with acyclovir, valacyclovir, or famciclovir — was associated with an adjusted hazard ratio of 0.092 for subsequent dementia. That is, a roughly tenfold reduction in risk.

This is a remarkable signal. It is also observational, and observational signals of this magnitude in pharmacoepidemiology almost always shrink when tested in randomised trials. The healthy-user effect is almost certainly part of the story: patients who receive and adhere to antiviral treatment tend to be more engaged with their health care, better educated, and less frail — all factors that independently predict lower dementia risk. The counterfactual — what would have happened to these patients had they not received antivirals — cannot be directly measured from cohort data alone.

A randomised trial is therefore essential. Several are now under way. The largest, currently enrolling, tests valacyclovir in carriers of the APOE $\epsilon 4$ allele — the strongest known genetic risk factor for late-onset Alzheimer's disease and a factor that may amplify HSV-1 reactivation in neurons. If that trial shows a positive effect, the infection hypothesis will have survived a genuine test; if it does not, Tzeng's 2018 finding will likely be regarded as a large confounded association rather than a causal effect. At the time of writing, the answer is unknown.

The revised Alzheimer's framework

One more change in the Alzheimer's landscape is worth flagging because it reshapes what “treating Alzheimer's” can even mean. In 2024 and 2025, the diagnostic criteria for Alzheimer's disease were revised to recognise the disease on the basis of biomarkers — amyloid PET positivity, or

cerebrospinal-fluid A β ₄₂/A β ₄₀ ratios, or specific blood-based phosphorylated-tau assays — even in cognitively normal individuals. A person who has no memory problems but has a positive amyloid scan now meets the biological criteria for Alzheimer’s disease. This change is controversial in clinical neurology for reasons that go well beyond our scope. But for the microbiome field it is significant: it extends the window in which disease-modifying interventions could plausibly act by perhaps a decade or more before symptoms would have appeared. If gut-directed interventions are to help at all, they are far more likely to help in biomarker-positive, cognitively normal individuals — the population the new criteria now formally identify — than in patients whose cortex is already visibly atrophied.

The clinical trial of the 2030s, if microbiome-based prevention of Alzheimer’s is to happen, will be a long, large, biomarker-enriched study of a microbiome-targeted intervention in the prodromal window. None of the trials currently under way look quite like that. Until one is run, microbiome-based prevention of Alzheimer’s remains an elegant idea with, at best, observational support.

29.3 Depression and the Promise of Microbiome-Targeted Interventions

Of the disease categories in this chapter, depression is the one most often discussed in the popular press as having a gut-based component. The phrase “gut feeling” is ubiquitous; the idea that serotonin is made in the gut is widely repeated, usually without the important caveat that gut-produced serotonin does not cross the blood–brain barrier and is a poor model for the central serotonin systems that antidepressants target. The chapter on the gut–brain axis, Chapter 13, built the careful version of this story. Here we extend it to the clinical questions: what have human trials shown, and what do they mean?

A brief recap from Chapter 13 is warranted. A 2019 paper from Mireia Valles-Colomer and colleagues, studying more than a thousand Flemish participants, reported that individuals with a clinical diagnosis of depression were consistently depleted of two butyrate-producing genera, *Coprococcus* and *Dialister*, after controlling for antidepressant use and other confounders [REF:vallescolomer2019]. John Cryan and Ted Dinan, working at University College Cork, had by then already articulated the gut–brain axis as a formal framework, with vagal afferents, short-chain fatty acids, tryptophan/kynurenine metabolism, and microbially produced GABA as the principal mechanistic channels. A 2016 experiment by John Kelly and colleagues showed that stool from patients with major depression, transferred into germ-free rats, produced changes in behaviour consistent with depression and in tryptophan metabolism consistent with altered kynurenine signalling [REF:kelly2016]. The basic framework — that human depression has microbial correlates and that transferring those correlates into a rodent can produce rodent analogues of depressive behaviour — was established before this book began.

What has happened since is the clinical era.

Psychobiotics and the meta-analytic verdict

The term “psychobiotic” was coined by Cryan and Dinan in 2013 to describe live organisms that, when ingested in adequate amounts, produce mental-health benefits. A decade of randomised controlled trials has asked whether psychobiotics work, and if so, which ones, in which patients.

The most recent large synthesis is the 2025 systematic review and meta-analysis by Asad and colleagues in *Nutrition Reviews*, which we introduced briefly in Chapter 13 [REF:asad2025]. Its scope

is psychobiotic interventions in patients with clinically diagnosed depression or anxiety, as distinct from participants in the general population reporting sub-clinical symptoms. Across the included trials, psychobiotic interventions produced a small but statistically significant reduction in depressive symptoms — with typical standardised mean differences favouring the intervention over placebo at magnitudes in the -0.3 to -0.5 range over four-to-twelve-week treatment periods. Anxiety symptoms showed similar, somewhat smaller, improvements.

Two features of this evidence base are worth dwelling on. The first is that the effect sizes, while real, are substantially smaller than those of standard antidepressants. An effective SSRI typically produces a standardised mean difference against placebo of around 0.3 in large randomised trials — itself a famously modest effect that has generated decades of argument. Psychobiotic effects are in the same ballpark. No trial to date has tested a psychobiotic as monotherapy against an SSRI, and there is no evidence that psychobiotics are a substitute for pharmacotherapy in moderate or severe depression. They are best understood, at present, as a plausible adjunct with modest incremental benefit.

The second feature is strain-specificity. The trials that have produced the strongest individual signals have tended to use specific, well-characterised strains — *Lactiplantibacillus plantarum* 299v in treatment-resistant depression, *Bifidobacterium longum* 1714 in stress and anxiety, *Bifidobacterium breve* CCFM1025 in subthreshold depression. Other strains of the same genera have shown little or no effect. A 2025 review in *Nutrients* of 45 trials confirmed that *Lactobacillus* species (broadly) appeared in about 45 per cent of trials and *Bifidobacterium* in about 29 per cent, but identified no specific cross-trial “best” strain [REF:huang2025]. This is the now-familiar strain-specificity problem from Chapter 23, and it has practical consequences. A consumer who buys a generic “probiotic for mood” is very unlikely to be buying the strain that performed best in trials. The category as it is sold bears only a loose relationship to the category as it is studied.

FMT for depression

The more radical intervention is not a single strain but the whole community. Over the past five years, a small set of randomised controlled trials has asked whether faecal microbiota transplantation — the nuclear option of Chapter 24 — can produce clinically meaningful improvements in patients with major depression.

The early trials, from 2021 to 2023, were pilot or open-label studies in Canada and Australia. They established feasibility and acceptability, showed that patients would tolerate FMT for mood indications, and produced heterogeneous efficacy signals that were too small to generalise from. A 2025 meta-analysis in *Frontiers in Psychiatry* pooled 12 RCTs involving 681 participants across China, Australia, Canada, Finland, and the United States [REF:depressionfmtmeta2025]. The standardised mean difference for depressive symptoms was -1.21 in favour of FMT, with a p-value of 0.0003 — a large effect on paper. Direct gastrointestinal delivery (nasoduodenal or colonoscopic) outperformed capsules. Benefits were stronger in patients with comorbid irritable bowel syndrome, a subgroup in which improvements in gut symptoms may have contributed to overall wellbeing. By six months, the effect had substantially attenuated.

A competing 2026 meta-analysis of seven studies with 235 participants reported no significant effect of FMT on depressive symptoms [REF:depressionfmtmeta2026]. The choice of which studies to include, and how to weight them, explains much of the divergence. The honest summary is that the evidence base is inconsistent, the effect sizes vary dramatically across trials, and the durability of any effect is questionable. FMT for depression is, at the time of writing, at the stage FMT for

Clostridioides difficile infection was in 2008 — plausible, safe, suggestive, but not ready for clinical practice. The analogy is not quite exact. *C. difficile* infection is an obvious dysbiosis disease with a clear ecological logic for FMT as replacement therapy. Depression's relationship to the microbiome is far less direct, and the mechanistic rationale for replacing the whole community is correspondingly weaker.

The inflammatory biotype

Why, if the mouse data are strong and the mechanistic story plausible, have the human trials produced such inconsistent results? One compelling hypothesis — advanced most forcefully by Carmine Pariante and colleagues at the Institute of Psychiatry in London, but widely shared across the field — is that major depressive disorder is not a single condition but a mix of biologically distinct subtypes, only one of which is likely to respond to microbiome-targeted interventions [REF:miller2020].

The inflammatory-subtype hypothesis of depression proposes that roughly 30 per cent of patients with major depression carry a detectable signature of chronic, low-grade systemic inflammation — elevated CRP and IL-6, raised white blood cell counts, alterations in cytokine profiles. These patients respond poorly to standard SSRIs and may respond better to anti-inflammatory strategies. The gut microbiome, with its ability to drive or dampen systemic inflammation through the barrier-and-LPS pathway we discussed in Section 29.2, is the most biologically plausible upstream driver of this inflammatory subtype. If that is right, the patients most likely to respond to FMT or to anti-inflammatory psychobiotics are exactly the patients who usually do worst on conventional antidepressants — a clinically important mismatch that standard trial designs, which enrol across the spectrum of depressive illness, systematically miss. Trials that enrolled all-comers would be expected to show small, heterogeneous effects, because the true responders are diluted by larger numbers of non-responders. This is, plausibly, the pattern the current literature shows.

If the inflammatory-subtype model holds up — and the evidence is suggestive rather than definitive — the path forward is a trial that enrolls only inflammation-positive, treatment-resistant patients and randomises them to FMT versus sham procedure with long-term follow-up. Such a trial has not yet been done. It is, in our view, the single most informative experiment the depression-microbiome field could run.

The vagus nerve is a complementary part of this story. Vagal afferents carry visceral inflammatory and metabolic signals from gut and other peripheral organs to the brainstem, where they influence mood circuits in ways that are now fairly well mapped. Vagus nerve stimulation — a surgical implant that delivers electrical pulses to the vagus in the neck — is approved for treatment-resistant depression in several jurisdictions. Its efficacy is modest but real. The mechanistic substrate it engages is presumably the same pathway that psychobiotics and FMT would need to engage to produce durable mood effects. One reading of the current evidence, then, is that vagus nerve stimulation and psychobiotic interventions are two ends of the same therapeutic rationale, with the former delivering a reliable but invasive electrical signal and the latter attempting, much more diffusely, to achieve the same through ecological modulation.

The edges: bipolar disorder, schizophrenia, and the limits of generalisation

The gut-brain axis framework has, predictably, been extended to other psychiatric conditions, with generally weaker signals. Bipolar disorder and schizophrenia both show gut microbiome differences from healthy controls in case-control studies, but both are heavily confounded by the effects

of psychotropic medication on gut ecology. The atypical antipsychotics, in particular, produce profound microbiome disturbances through their known effects on weight, glucose metabolism, and gastrointestinal motility. Disentangling disease-related from treatment-related dysbiosis is methodologically difficult and, in many of the published studies, not persuasively attempted. The careful reader of the psychiatric microbiome literature should, at the moment, reserve judgment.

Interlude: Autism Spectrum Disorder

A chapter on the microbiome and the brain that did not discuss autism spectrum disorder would be incomplete, because the ASD–microbiome literature is one of the most active and one of the most controversial in the field. It is also partly covered elsewhere in this book. In Chapter 14 we met the elegant experiments of Mauro Costa-Mattioli’s group showing that a specific *Lactobacillus reuteri* strain could rescue social behaviour deficits in a mouse model of ASD by signalling through vagal afferents and altering oxytocin release [REF:buffington2016; REF:sgritta2018]. That work remains among the strongest mechanistic demonstrations of a gut–brain behavioural axis in any animal model. Here we extend the clinical picture.

The best-known clinical evidence for microbiome-based intervention in ASD comes from Dae-Wook Kang and James Adams’s 2017 *Microbiome* paper on “microbiota transfer therapy,” or MTT [REF:kang2017]. Kang and Adams enrolled 18 children, aged 7 to 16, with ASD and chronic gastrointestinal symptoms. Participants received a two-week course of oral vancomycin to suppress the existing microbiota, a bowel cleanse to clear residual colonic contents, a large initial dose of standardised donor stool administered orally and rectally, and then seven to eight weeks of maintenance oral FMT. The results were striking. Gastrointestinal Symptom Rating Scale scores fell by around 80 per cent from baseline. Behavioural measures of ASD severity — the Childhood Autism Rating Scale, the Social Responsiveness Scale, parent-rated problem behaviours — improved substantially. Microbial analyses showed increases in *Bifidobacterium* and *Prevotella* that persisted after the course ended.

Kang and Adams followed their participants for two years and published the long-term results in 2019 [REF:kang2019]. Most of the GI improvements were maintained. Behavioural scores had continued to improve, in some measures reaching the non-ASD range. The microbial changes persisted. The study was striking enough to generate enormous interest from families and from the field.

It was also, methodologically, an open-label trial with no placebo arm, no blinding, and a sample size of 18. The baseline was a severely affected clinical population whose parents had chosen to enrol in a demanding multi-week protocol. Parents administering a demanding protocol for their child are not unbiased raters of behavioural change. Regression to the mean and expectancy effects are substantial in ASD because the phenotype fluctuates substantially month-to-month in many children. The MTT results are compelling enough to have generated a follow-on randomised controlled trial, which at the time of writing is in progress and has not yet reported. Until it does, the MTT evidence base remains a single open-label trial of 18 children — an excellent pilot, but not, in the strict evidence-based sense, a proven intervention.

On the mouse side, a 2019 study from Gil Sharon and colleagues in the Mazmanian laboratory extended the transfer paradigm to ASD [REF:sharon2019]. Germ-free mice colonised with stool from human donors with ASD showed more mouse-autism-like behaviour than mice colonised

with stool from neurotypical donors. Brain tissue from the ASD-colonised mice showed alternative splicing of genes previously implicated in autism, and specific microbial metabolites — notably 5-aminovalerate and taurine — were sufficient, when given alone, to improve behaviour in a separate mouse model of ASD. This is the transfer experiment that, in humans, is not permissible; in mice, it provides mechanistic support for a gut–brain contribution to ASD-like behaviour.

A particular metabolite, 4-ethylphenylsulfate (4-EPS), has become the best-characterised candidate mediator in this literature. Elaine Hsiao, working with Sarkis Mazmanian, showed in 2013 that the maternal-immune-activation mouse model of ASD-like behaviour has plasma 4-EPS levels elevated roughly 46-fold over controls, that treatment with *Bacteroides fragilis* normalises 4-EPS levels and improves behaviour, and that administration of 4-EPS alone to naive mice produces anxiety-like behaviour [REF:hsiao2013]. A 2022 *Nature* paper from Brittany Needham and colleagues extended the mechanism: 4-EPS appears to cross the blood–brain barrier and alter oligodendrocyte maturation and myelination in behaviourally relevant circuits [REF:needham2022]. Another metabolite, p-cresol, has been shown in a 2021 study from Nicolas Glaichenhaus’s group to induce ASD-like behaviours in mice, to be elevated in many children with ASD, and to act through reduced activity of dopamine neurons in the social reward circuit [REF:bermudezmartin2021].

This is among the strongest mechanistic literature for any gut–brain behavioural link. It is also — and this is the caveat that deserves emphasis — a literature built on mouse models of human ASD. “Mouse autism-like behaviour” is operationally defined as time spent sniffing an unfamiliar conspecific, time spent burying marbles, time spent grooming, or time spent in the centre versus periphery of an open arena. These are not diagnostic proxies for human ASD. Human autism is a neurodevelopmental condition defined by differences in social communication, by sensory processing differences, by restricted or repetitive patterns of behaviour, and by a range of intellectual and adaptive profiles that no mouse assay captures. The mouse findings tell us that specific microbial metabolites can change specific behavioural outputs in specific mouse models. They do not tell us that ASD is a gut disease, that probiotics cure autism, or — and this is a warning to any parent reading — that FMT is an appropriate treatment for their autistic child. The evidence does not yet support any of those conclusions, and the ASD community has been on the receiving end of enough premature therapeutic claims over the decades to deserve a careful, honest evaluation here.

The verdict, then: the mouse evidence for a gut–brain contribution to ASD-like behaviour is the strongest of any neuropsychiatric condition covered in this chapter. The human evidence — one open-label trial of 18 children, with a randomised follow-up pending — is not yet sufficient to change practice. Families seeking intervention should know this carefully.

The edge of the map: multiple sclerosis and ALS

Two further neurological conditions — and one surprising peripheral story — are worth flagging before we turn to the question of translation. The first two are central nervous system diseases where the microbiome is implicated; the third is a peripheral neuropathy that may, in some cases, originate not in a pill bottle but in the gut.

Multiple sclerosis is partly an autoimmune disease, and Chapter 26 treated the MS–microbiome literature in detail. The landmark papers are there: Jangi’s 2016 characterisation of the MS gut microbiome [REF:jangi2016], the elegant monozygotic-twin experiments from Kerstin Berer and Gurumoorthy Krishnamoorthy showing that stool from the MS-affected twin, transferred into a

transgenic mouse model of demyelinating disease, produced more autoimmune pathology than stool from the unaffected twin [REF:berer2017], and Veit Rothhammer’s demonstration that microbial tryptophan metabolites — specifically indoles — act on the aryl hydrocarbon receptor (AHR) on astrocytes to dampen central nervous system inflammation [REF:rothhammer2016]. The punchline of the MS story, covered in Chapter 26, is that gut microbes produce metabolites that actively modulate the central nervous system’s immunological setpoint, and that disturbances of the microbiota in MS patients may tilt that setpoint toward disease. We note here only that MS is the condition in which the causal arrow between gut and brain is most clearly established in mouse models, and that the same pathway — microbial tryptophan metabolism through AHR — has now been implicated in the broader biology of neurodegeneration. We refer the reader to Chapter 26 for the fuller treatment.

Amyotrophic lateral sclerosis (ALS) — the rapidly progressive motor neuron disease that killed Stephen Hawking much later in life than is typical — is a harder and newer story. The seminal contribution came from Eran Blacher, working in Eran Elinav’s laboratory, in 2019 [REF:blacher2019]. Blacher studied a transgenic mouse model of ALS carrying a mutation in the SOD1 gene (a small minority of human ALS is caused by SOD1 mutations, but the mouse is a tractable laboratory analogue). SOD1-mutant mice show gut dysbiosis before they show motor symptoms — another prodromal signature. Supplementation with *Akkermansia muciniphila* ameliorated motor disease; supplementation with *Ruminococcus torques* or *Parabacteroides distasonis* worsened it. *A. muciniphila*-associated nicotinamide accumulated in the central nervous system; systemic nicotinamide itself improved motor symptoms in affected mice. A small human cohort showed altered microbial signatures in ALS patients and differences in serum nicotinamide concentration that paralleled the mouse findings, though the human data are noisier.

The twist — and it is a twist that deserves attention — came in subsequent work on a different ALS model. Mice carrying a mutation in the C9orf72 gene (the other common genetic cause of human ALS) show the *opposite* microbiome effect: depletion of the microbiota worsens the disease in SOD1 mice but ameliorates it in C9orf72 mice. The same intervention, in the same disease, in the same species, produces opposite effects depending on the underlying genetic cause. This is a vivid reminder that there is no such thing as an “ALS microbiome” — there are microbiome effects that interact with specific genetic backgrounds. Much of the human ALS literature has not been powered to detect this kind of subtype-specific signal, and future progress will likely require genotype-stratified clinical studies. For now, the ALS story is a reminder that the gut–brain axis is not a single channel but a bundle of interacting ones, and that unifying narratives about “the microbiome of neurodegeneration” should be treated with the same caution we applied to “the cancer microbiome” in Chapter 28.

The Vitamin Factory Gone Wrong: Microbiome-Produced B6 and Peripheral Neuropathy

The diseases discussed so far in this chapter — Parkinson’s, Alzheimer’s, depression, autism, multiple sclerosis, ALS — are all diseases of the central nervous system. The microbiome reaches them indirectly, through vagal signalling, through inflammatory mediators that cross the blood–brain barrier, through metabolites that alter central immune cells. The barrier itself, as we have emphasised, is a formidable obstacle. But there is a part of the nervous system that sits *outside* the blood–brain barrier, directly exposed to whatever the bloodstream carries — and it is precisely this

anatomical vulnerability that makes the next story possible.

The dorsal root ganglia (DRG) are clusters of sensory neuron cell bodies that sit along the spinal column, just outside the spinal cord. They are the relay stations for touch, temperature, pain, and proprioception — the signals that tell you where your limbs are in space and whether the bathwater is too hot. Unlike motor neurons, whose cell bodies are tucked safely inside the spinal cord behind the blood–brain barrier, the DRG neurons are bathed in a capillary bed that is far more permeable. They are exposed to circulating metabolites, drugs, and toxins at concentrations that the brain and spinal cord never see. This is why certain chemotherapy agents cause peripheral neuropathy while sparing the brain: the drugs reach the DRG but cannot cross the blood–brain barrier. And it is why vitamin B6, in excess, targets the same neurons.

Vitamin B6 toxicity is one of the oldest and most reliably reproduced forms of nutritional neurotoxicity. The clinical syndrome — a sensory neuropathy characterised by numbness, tingling, burning, and loss of proprioception in the hands and feet, progressing proximally — was first described in detail in the 1980s. Schaumburg and colleagues in 1983 published the landmark human case series; Xu and colleagues in 1989 showed that high-dose pyridoxine in rats caused dose-dependent necrosis of DRG sensory neurons, with the largest-diameter neurons — those with the longest axons and greatest metabolic demands — dying first [REF:xu1989]. Hoover and Carlton demonstrated the same lesion pattern in beagle dogs: degeneration confined precisely to tracts and nerves whose cell bodies sit in the spinal and trigeminal ganglia — the neurons outside the barrier [REF:hoover1981]. The clinical picture is now well established and forms part of the standard neurological differential for sensory-predominant peripheral neuropathy [REF:alvarez2026; REF:snodgrass1992].

What makes the B6 story remarkable — and what connects it to this book — is a paradox.

The Vitamin B6 Paradox

Vitamin B6 is not a single molecule. It is a family of six related compounds called **vitamers**: pyridoxine (PN), pyridoxal (PL), pyridoxamine (PM), and their phosphorylated forms — pyridoxine 5'-phosphate (PNP), pyridoxal 5'-phosphate (PLP), and pyridoxamine 5'-phosphate (PMP). Of these six, only one is biologically active: **pyridoxal 5'-phosphate (PLP)**. PLP is one of the most versatile coenzymes in human biochemistry — it participates in over 150 enzymatic reactions, including the synthesis of neurotransmitters (serotonin, dopamine, GABA, noradrenaline), the metabolism of amino acids, the production of haemoglobin, and the regulation of gene expression. When clinicians measure “vitamin B6 status,” they are measuring PLP in plasma. Adequate status is generally defined as a PLP concentration above 30 nmol/L.

The supplements that people buy, however, are overwhelmingly **pyridoxine** — the inactive vitamer. To become useful, pyridoxine must be converted to PLP in a two-step process: first phosphorylated by the enzyme **pyridoxal kinase (PDXK)** to pyridoxine 5'-phosphate, then oxidised by **pyridoxine 5'-phosphate oxidase** to PLP. This works well at physiological doses. At high doses, it goes wrong.

In 2017, Misha Vrolijk and colleagues at Maastricht University published a study they titled “The vitamin B6 paradox” [REF:vrolijk2017]. Testing the six B6 vitamers individually on neuronal cell lines, they found that only one was neurotoxic: pyridoxine. The other five — including the biologically active PLP — caused no cell death at any concentration tested. The paradox runs deeper: pyridoxine does not merely poison neurons through some generic toxic mechanism. It **competit-**

ively inhibits PLP. Excess pyridoxine floods the PDXK enzyme, saturating it and simultaneously blocking PLP-dependent enzymes from accessing their cofactor. The result is that taking too much vitamin B6 in supplement form produces a state biochemically indistinguishable from vitamin B6 deficiency — because the active form, PLP, is displaced from the enzymes that need it.

Hadtstein and Vrolijk extended this into a comprehensive mechanistic model in 2021 [REF:hadtstein2021]. Their proposed pathway runs as follows. Excess circulating pyridoxine inhibits PDXK in peripheral tissues. PDXK inhibition reduces local PLP availability. Reduced PLP availability impairs the enzyme **glutamic acid decarboxylase (GAD65)**, which requires PLP to synthesise **GABA** — the principal inhibitory neurotransmitter. In sensory neurons of the DRG, GABA signalling modulates action potential filtering at the T-junction — the point where the peripheral and central branches of a sensory axon meet. Without adequate GABA, the filtering breaks down, leading to aberrant depolarisation, excessive calcium influx, excitotoxicity, and ultimately neuronal death.

The reason the brain is spared is now clear: pyridoxine is relatively impermeable to the blood–brain barrier. In the brain, PDXK activity is preserved and PLP-dependent enzymes continue to function. Motor neurons are spared because their cell bodies sit within the spinal cord, also behind the barrier. But the DRG, sitting outside the barrier, are exposed to full circulating pyridoxine concentrations. They are the first casualties — and, in many patients, the only ones.

Where the Microbiome Enters

Until recently, vitamin B6 toxicity was understood as a disease of supplementation — a consequence of megadose pyridoxine pills, typically at intakes above 50 mg per day sustained over months, and most reliably at intakes above 200 mg per day. But clinicians have occasionally encountered patients with elevated serum B6 and peripheral neuropathy whose reported supplement intake seemed far too low to explain the levels. A 2021 review by Reddy noted that the association between elevated pyridoxine and neuropathy is not always traceable to obvious exogenous sources, and that B6 metabolites have long half-lives that can accumulate unpredictably [REF:reddy2021]. This leaves a gap in the clinical story: where is the extra B6 coming from?

One emerging answer is the gut microbiome.

As we noted in Chapter 15 — and as the reader will recall from the broader theme of microbial metabolism running through this book — the gut microbiome is a prolific vitamin factory. Metagenomic analyses estimate that microbiota-derived vitamins could contribute 27 to 86 per cent of reference intakes for several B vitamins, including B6. The bacteria that possess de novo B6 biosynthesis pathways include *Bacteroides fragilis*, *Prevotella copri*, *Bifidobacterium longum*, and *Collinsella aerofaciens* — all common residents of the healthy human colon [REF:magnusdottir2015]. A large population study by Radjabzadeh and colleagues found that pyridoxine biosynthesis pathways were significantly overrepresented in the gut microbiomes of children compared with adults — suggesting that microbial B6 production varies substantially with community composition and age [REF:radjabzadeh2020].

Crucially, the vitamin that gut bacteria tend to produce is **pyridoxine** — the inactive form, and the one associated with neurotoxicity. Wu and colleagues in 2025 demonstrated that *Lactobacillus fermentum* and *Pediococcus acidilactici*, isolated from gut environments, synergistically synthesise vitamin B6 de novo, maintaining it primarily as pyridoxine (84 per cent of their output) and PMP (16 per cent) [REF:wu2025]. This is not an incidental detail. If the microbiome's B6 output is overwhelmingly pyridoxine rather than PLP, then microbiome-derived B6 is, at least in principle,

subject to the same paradox as supplemental pyridoxine: at high enough concentrations, it could inhibit PLP-dependent enzymes and produce the functional deficiency that leads to neuropathy.

The question, then, is whether the microbiome can produce enough pyridoxine to matter.

The Dysbiosis Amplifier

Under normal circumstances, the answer is probably no. The colonic microbiome produces B6, but the colon is not the primary site of B6 absorption — the jejunum is. B6 synthesised in the colon may contribute modestly to systemic levels through passive absorption, but the quantitative contribution under healthy conditions is unlikely to reach neurotoxic concentrations.

Dysbiosis, however, can change the equation.

Small intestinal bacterial overgrowth — SIBO — shifts the site of microbial metabolism upstream, into the jejunum and ileum, where absorption is maximal. If B6-producing bacteria overpopulate the small intestine, pyridoxine synthesised there would be absorbed far more efficiently than pyridoxine produced in the colon. Clinical commentary has proposed that methane-dominant SIBO, in particular, may drive elevated B6 levels through exactly this mechanism — and that patients with unexplained elevated B6 and peripheral neuropathy should be evaluated for SIBO as a potential contributing cause.

Broader dysbiosis may also contribute. Abdollahiyan and colleagues, studying gut microbiomes in patients with non-alcoholic fatty liver disease, found that the **pyridoxal 5'-phosphate biosynthesis pathway** was significantly more abundant in NAFLD patients than in healthy controls [REF:abdollahiyan2023]. A 2024 study of chronic stress in animal models showed that stress-induced dysbiosis altered vitamin B6 metabolite levels in plasma — including 4-pyridoxic acid, the principal urinary excretion metabolite of B6 — in ways that correlated with specific genus-level changes in the gut community [REF:gut_dysbiosis_b6_2024]. The microbiome's vitamin B6 output, it appears, is not fixed. It shifts with the community, and disturbances to the community can shift it meaningfully.

An Unproven but Biologically Compelling Hypothesis

It is important to be precise about what the evidence does and does not support. Each link in the chain is individually well established: bacteria synthesise B6, predominantly as pyridoxine; excess pyridoxine causes peripheral neuropathy through competitive inhibition of PLP; dysbiosis can alter microbial vitamin production; SIBO shifts production upstream to where absorption is maximal. What does **not** yet exist is a published human study demonstrating the complete causal chain — from a specific dysbiotic state, through elevated systemic pyridoxine, to clinical neuropathy — in a patient who is not taking B6 supplements. The hypothesis is biologically compelling. It has not been proven.

This matters clinically for two reasons. First, it suggests that patients with unexplained peripheral neuropathy and elevated serum B6 — particularly those who deny supplement use — should have their gut health evaluated, including testing for SIBO. If microbial overproduction of pyridoxine is contributing to their B6 burden, treating the dysbiosis may be as important as stopping supplements. Second, it adds a layer of nuance to the familiar narrative of the microbiome as benign provider of essential nutrients. The microbiome's vitamin factory, like any factory, can overproduce — and the consequence of overproduction, in this case, is not waste but nerve damage.

There is also a broader lesson here about the nature of microbial metabolites. Throughout this book we have emphasised that the same molecule can be beneficial or harmful depending on its concentration, its location in the gut, and the physiological context of the host. Short-chain fatty acids are essential at normal levels and problematic in excess. Bile acids regulate metabolism when properly balanced and promote disease when the balance is disturbed. Vitamin B6 — the molecule whose active form participates in over 150 biochemical reactions — joins this list. The microbiome does not produce “good” or “bad” metabolites. It produces metabolites whose effects are dose-dependent, context-dependent, and occasionally, when the ecology shifts, capable of causing the very deficiency they were supposed to prevent.

29.4 The Translation Gap: From Mouse Models to Human Trials

We have come, in four passes through four disease stories, to a recurring observation. In every case the pattern is broadly similar: exciting mouse findings, supportive human observational data, small and methodologically modest clinical trials with mixed or modest effects, and a large gap between the press coverage and what the evidence strictly supports. The temptation, for a reader who has followed all of this, is to dismiss the field as overhyped. That would be wrong. The temptation, for a reader who wants to believe, is to let the most exciting findings stand unchallenged. That would also be wrong.

This section is a deliberate pause in the disease narrative. Its purpose is to teach the reader how to read the next microbiome-and-brain paper they encounter — whether it is a *Cell* study on a new gut–brain metabolite, a *Nature* paper on a germ-free mouse model of some condition, or a press release announcing positive results from a small clinical trial. The reading skills this section tries to develop are not specific to microbiome science. They are the general skills of calibrated scepticism that apply to any young biomedical field in its formative decade.

Germ-free mice are not empty humans

The single most important technical tool of the mouse-era gut–brain literature is the germ-free mouse: an animal born and raised in a sterile isolator, exposed from the moment of Caesarean-section birth to nothing but filtered air, autoclaved food, and sterile bedding, and thereby possessing no microbiota of any kind. Germ-free mice can be studied as they are, or colonised — mono-associated with a single species, di-associated with a defined pair, or gavaged with a full community from another animal or a human — to ask whether a given microbial community is sufficient to produce a given phenotype. This is, to a first approximation, the closest thing the field has to a causal experiment in a whole vertebrate organism.

It is also, as every gnotobiotic animal facility quietly knows, a deeply imperfect tool. Germ-free mice are not mice with a missing microbiome. They are mice whose entire development has proceeded in the absence of microbial signals, and they are structurally abnormal as a consequence. Their gut-associated lymphoid tissue is rudimentary. Their Peyer’s patches are small and their germinal centres under-developed. Their intestinal mucus layer is thinner than conventional mice. Their intestinal epithelium has abnormal transit and proliferation kinetics. Their blood–brain barrier is more permeable than that of conventional mice. Their hypothalamic–pituitary–adrenal axis responds more vigorously to stressors. Their hippocampus has less brain-derived neurotrophic factor. Their prefrontal cortex shows abnormal myelination. Their behaviour is abnormal in nearly every assay that has been tried — more anxiety-like, more hyperactive, differently social.

A clear-eyed methodological review by Maaïke Jans and colleagues in 2025 set out these limitations carefully [REF:jans2025]. Germ-free mice are an excellent tool for answering one question: is the phenotype of interest microbiome-dependent at all? If the germ-free mouse lacks the phenotype and the conventional mouse has it, we have learned something important. But if the addition of a specific bacterium to a germ-free mouse restores a phenotype, we have not learned that this bacterium causes the phenotype in a normal animal. We have learned that this bacterium is sufficient, in a developmentally abnormal host, to rescue a phenotype whose normal baseline is itself abnormal. These are very different claims.

The same caution applies, with even more force, to “human microbiota–associated” or HMA mice — germ-free mice colonised with stool from a human donor. HMA mice are widely used to test causal hypotheses: give the mice stool from a human patient, ask whether the mouse develops the relevant phenotype, and conclude that the human microbiome is sufficient for the phenotype. The methodological problem is that HMA mice are poor ecological containers for human microbes. The human microbiota is adapted to the human gut’s particular pH, bile acid repertoire, mucin composition, and temperature; the mouse gut has none of these in the same configuration. What tends to happen in practice is that HMA mice, regardless of donor, converge toward a mouse-like community dominated by a small number of human-derived species that happen to tolerate the mouse gut. Different recipients of the same donor do not always look alike; the same donor instilled on different days does not always take in the same way. Conclusions about “the human microbiome causes X” based on HMA-mouse phenotypes are built on an ecological substrate that is nothing like the original human community. The conclusions may still be right, sometimes. But they are almost never as clean as the papers imply.

What mouse behavioural endpoints actually measure

The second methodological issue, and perhaps the more important one for neuropsychiatric translation, is what mouse behavioural tests actually measure. The canonical assays — developed mostly in the 1970s and 1980s for antidepressant drug screening in the pharmaceutical industry — have names that sound like they measure human symptoms. They do not.

The open field test places a mouse in a brightly lit rectangular arena and measures how much time it spends in the centre versus the periphery, on the reasonable premise that more anxious mice stay near the walls (thigmotaxis) and less anxious mice venture into the open. The open field test measures locomotion and thigmotaxis. It does not measure anxiety as a human being experiences it — the cognitive rehearsal of catastrophe, the somatic dread, the avoidance behaviour structured around specific feared situations. A drug that increases a mouse’s centre-time may make it a less anxious mouse or may simply make it a more active one; the two are difficult to dissociate without careful controls.

The forced swim test drops a mouse into a cylinder of water it cannot escape and measures how long it swims before giving up and floating. Less swimming — less “struggle” — was originally interpreted as a model of depression-like “despair,” and antidepressant drugs were shown to increase struggle time. The test is now widely criticised as a poor model of depression. It probably measures learned helplessness or energy economy, not the clinical phenomenon of major depressive disorder as experienced by human patients. It remains in use mostly because an enormous existing dataset of drug responses exists in its coordinate system, which is a different argument from saying it is valid.

The three-chamber social preference test places a mouse in a box with three compartments, one of

which contains an unfamiliar conspecific, and measures which compartment the test mouse prefers to spend time in. It is widely used as a model of “social behaviour” in autism research. It measures novelty preference and exploratory behaviour more than it measures anything we would recognise as social connection. A drug that makes a mouse approach a novel conspecific more readily may be changing its social motivation or may simply be increasing its exploratory drive toward any novel stimulus.

Marble-burying — a test in which a mouse is placed in a cage with glass marbles and the number of marbles buried under bedding is counted — is sometimes used as a model of “repetitive behaviour” in autism research. It measures how much digging a particular mouse does. Mice dig in response to a wide range of conditions, not all of which map obviously onto human repetitive behaviours.

We dwell on these tests because the microbiome-and-brain literature is saturated with claims like “strain X reduces anxiety-like behaviour in mice” or “microbial metabolite Y improves social behaviour in a mouse model of autism.” Every such claim is a claim about what one or more of these tests returns. A probiotic that improves a mouse’s open-field performance may have done many things other than reduce its “anxiety”; increased locomotion in an unfamiliar environment is only one of many possibilities, and it is the least interesting for clinical translation. The reader who cannot evaluate the endpoint cannot evaluate the claim. We have tried to signal, throughout this chapter, when a cited study’s endpoint is a mouse behavioural test versus a human clinical outcome, because those are very different kinds of evidence and should carry different weight.

The Cryan and Mazmanian framework

In 2022, John Cryan and Sarkis Mazmanian published a short framework paper in *Science* that attempted to synthesise what the microbiome-and-brain field would need to produce, over the coming decade, to earn clinical credibility [REF:cryan2022]. Its argument is worth paraphrasing in full.

Cryan and Mazmanian argued that association studies — cross-sectional comparisons of disease and healthy microbiomes — have taken the field as far as they can. What is needed next, they argued, are five complementary lines of work. First, longitudinal cohort studies that capture microbiome samples before clinical symptoms appear, and that therefore have the chance to show prospectively that microbial signatures precede disease. These studies are expensive and slow, and they are the reason the field’s strongest causal inferences about Parkinson’s disease (from RBD cohorts) and preclinical Alzheimer’s (from biomarker-ascertained populations) have been worth the decade-long effort. Second, Mendelian randomisation using microbial quantitative trait loci — human genetic variation that influences microbiome composition — to ask whether genetically driven microbial differences causally influence disease risk, in the same way Mendelian randomisation has been used to adjudicate the cholesterol-cardiovascular-disease causal story. Third, intervention trials of defined microbial consortia or specific microbial metabolites, rather than whole-stool FMT, because defined interventions are easier to standardise, reproduce, and regulate than the ecological black box of another human’s stool. Fourth, mechanistic dissection in gnotobiotic animals using defined communities of tractable size, rather than the whole human microbiota in a mouse gut. Fifth, cross-species triangulation — consistent findings across mouse, non-human primate, and human evidence — rather than over-reliance on any single model.

Crucially, Cryan and Mazmanian argued that prospective targeting of the gut-to-brain axis is feasible and timely even while fundamental causality gaps remain. Their point is not that the field should wait until every mechanism is proven before running trials, but that it should design trials that teach us something whether they succeed or fail — trials with defined interventions, pro-

spectively selected populations, and biomarker-backed endpoints that illuminate the mechanism regardless of the clinical outcome. Almost none of the trials cited in this chapter meet that bar. GUT-PARFECT comes closest; GAIN, for all its rigour, used a whole-population enrolment that could not stratify on the biology most likely to matter.

Hill's criteria, adapted

Austin Bradford Hill, in 1965, articulated a set of criteria for judging whether an observed association between an exposure and an outcome was plausibly causal. Hill's criteria — strength, consistency, specificity, temporality, dose-response, plausibility, coherence, experiment, and analogy — were designed for exposures like cigarette smoking and cancer, where the exposure and the outcome are both relatively discrete. They map imperfectly onto microbiome biology, where the “exposure” is a complex, shifting community rather than a single agent. But the mapping is informative.

Strength: most Parkinson's, Alzheimer's, and depression gut-microbiome signals have effect sizes that would qualify as moderate rather than strong — 1.5- to 3-fold differences between cases and controls in the abundance of specific taxa, with the Scheperjans Prevotella finding of 77 per cent depletion being unusually large. This is not a *Helicobacter*-and-gastric-cancer story; the effects are smaller and less specific.

Consistency: variable, often weak. The Romano 2025 meta-analysis showing cross-cohort AUC of only about 61 per cent for Parkinson's classifiers is a case in point. Findings do not port well across populations, and the “Parkinson's microbiome signature” exists, but imperfectly and not universally. The same problem appeared in the cancer-microbiome field and is being addressed by the microbial-fingerprint methodology we covered in Chapter 28; whether that approach will rescue neurological generalisability remains to be seen.

Specificity: poor. Gut dysbiosis is a feature of almost every chronic disease that has been examined. Butyrate producer depletion is a non-specific marker of systemic ill-health rather than a specific aetiological contributor to Parkinson's or Alzheimer's or depression. Whatever is driving disease-specific neuropathology, it is unlikely to be “low butyrate” on its own.

Temporality: increasingly strong. Prodromal-phase studies in RBD cohorts (for Parkinson's), in biomarker-positive but cognitively normal individuals (for Alzheimer's), and in preclinical ASD cohorts are now showing microbial signatures that precede clinical disease by years. This is the single strongest line of causal evidence in the field, because it addresses the most obvious alternative explanation — that dysbiosis is merely a consequence of being ill.

Dose-response: largely untested. The question of whether more extreme microbial disturbances produce more severe neurological outcomes has not been systematically examined, partly because “more extreme” is hard to define in an ecological framework.

Plausibility: strong. Well-characterised pathways exist through which gut microbes influence the brain: vagal afferents, short-chain fatty acids, tryptophan metabolism to indoles and kynurenines, bile acid signalling, LPS-driven inflammation, modulation of the blood–brain barrier. The biology is not exotic; it is standard physiology.

Coherence: the microbiome-and-brain hypotheses fit comfortably within wider knowledge of immune-brain interactions, developmental neuroscience, and neuroinflammation. They do not require the reader to discard any accepted biology.

Experiment: available in mice, only sketchy in humans. The FMT trials in Parkinson's and depression are the first real interventional experiments in humans; they are small and early.

Analogy: yes. The *H. pylori*–gastric cancer story, and the HPV–cervical cancer story, establish that microbial contributions to chronic diseases of specific organs are biologically plausible and clinically relevant. The Parkinson's and Alzheimer's microbiome stories are not stranger than those precedents.

Overall, the Hill's-criteria verdict for the microbiome's role in major neuropsychiatric disease is: temporality strong and gathering, plausibility strong, analogy strong, experiment tantalising but early, strength moderate, consistency patchy, specificity poor. That is the profile of a field in its adolescence — too advanced to dismiss, not yet mature enough to guide practice.

The failing and failed trials

A brief recapitulation. In a chapter that has covered the evidence in considerable detail, it is worth collecting in one place the randomised controlled trials of microbiome-targeted interventions in neuropsychiatric disease that have reported to date. GAIN (atuzaginstat for Alzheimer's): failed primary endpoint, suggestive subgroup signal in *P. gingivalis*–positive participants, programme discontinued. Multiple psychobiotic RCTs for depression: small effects, strain-specific, heterogeneous across populations. Xue 2023 (FMT for Parkinson's): positive on autonomic symptoms, modest motor effect, single-centre, small. GUT-PARFECT (FMT for Parkinson's): positive 3.1-point motor improvement at 12 months, single-centre, 46 participants, surprisingly durable. Kang MTT (FMT for ASD): open-label, 18 children, randomised follow-up pending.

What the field has not yet produced is a single large, multi-centre, double-blind, placebo-controlled, mechanistically informed randomised trial of a microbiome-targeted intervention with a primary clinical endpoint in neuropsychiatric disease. Until one exists, the appropriate position for the clinical community is: watchful interest, patient support of careful trials, and a refusal to let enthusiasm outrun evidence.

What would need to be shown

The trials that would move the field are visible even if they have not yet been conducted. We list them here briefly because doing so makes concrete what “more evidence” means in practice. A multi-centre randomised controlled trial of FMT in body-first Parkinson's disease (stratified on RBD and autonomic features) with two-year motor and biomarker follow-up. A randomised placebo-controlled trial of psychobiotic intervention in the inflammatory subtype of major depression (stratified on baseline CRP or IL-6) with proper blinding and plausible effect size expectations. A randomised controlled trial of antiviral prophylaxis in APOE ε4 carriers with biomarker-positive preclinical Alzheimer's disease, with biomarker and cognitive endpoints over five years. A randomised placebo-controlled trial of MTT in children with ASD and co-morbid gastrointestinal symptoms, stratified on baseline microbial fingerprint. These trials are expensive, long, and methodologically demanding. They are the kind of trial the field needs, and the kind of trial that, by the end of the decade, should have begun to report. If they produce positive results, we will have the beginnings of a new neuropsychiatric therapeutic modality. If they produce null results, the field will have learned something almost as valuable: where the ceiling lies.

29.5 Closing: A Field in Adolescence

If the cancer chapter was about a field approaching maturity, this chapter has been about a field in adolescence — tall, growing fast, capable of real things, prone to overclaiming, not yet ready to be trusted with the car keys. The temptation in writing a chapter like this is to resolve the dissonance either by celebrating the most exciting findings or by dismissing the lot. We have tried to do neither. The Parkinson's story is real, the Alzheimer's story is plausible, the depression story is suggestive, the autism story is mechanistically illuminating and clinically premature. None of this is nothing. None of it is yet enough to change what a neurologist or psychiatrist does on a Tuesday morning.

What should change — and what, we hope, this chapter has helped along — is the way readers evaluate the next wave of claims. The mouse experiments will become more sophisticated; the human cohorts will become larger; the trials will multiply; the press releases will get bolder. Some of the exciting findings will replicate; many will not. Progress in this field will look, in retrospect, a lot less like a cascade of breakthroughs and a lot more like patient, careful accumulation of evidence — most of it modest, some of it surprising, all of it needing the scepticism and the humility that this chapter has tried to model.

The book's final chapters now turn to the future of microbiome medicine as a whole — to diagnostics, to engineered therapies, to the ethical and regulatory choices that will shape the next decade. The brain, for reasons this chapter has tried to make plain, will be one of the slower tributaries of that river. That is not a discredit. It is a recognition that the hardest organ is the hardest for a reason, and that asking microbes to reach into it is a project whose rewards, if they come, will come slowly.

References

[REF:appendectomy_meta2025] Collaborative authors (2025). Association between appendectomy and subsequent risk of Parkinson's disease: a systematic review and meta-analysis. *Frontiers in Neurology*. (Specific citation to be confirmed during final proofing.)

[REF:arabi2023] Arabi TZ, Alabdulqader AA, Sabbah BN, Ouban A (2023). Brain-inhabiting bacteria and neurodegenerative diseases: the "Brain Microbiome" theory. *Frontiers in Aging Neuroscience*. (Cross-referenced with Chapter 11.)

[REF:asad2025] Asad MR, et al. (2025). Efficacy of probiotic, prebiotic, and synbiotic interventions in clinically diagnosed depression and anxiety: a systematic review and meta-analysis. *Nutrition Reviews*. (Cross-referenced with Chapter 13.)

[REF:bermudezmartin2021] Bermudez-Martin P, Becker JAJ, Caramello N, Fernandez SP, Costa-Campos R, Canaguier J, et al. (2021). The microbial metabolite p-Cresol induces autistic-like behaviors in mice by remodeling the gut microbiota. *Microbiome* 9(1):157. doi:10.1186/s40168-021-01103-z

[REF:berer2017] Berer K, Gerdes LA, Cekanaviciute E, Jia X, Xiao L, Xia Z, et al. (2017). Gut microbiota from multiple sclerosis patients enables spontaneous autoimmune encephalomyelitis in mice. *Proceedings of the National Academy of Sciences USA* 114(40):10719–10724. doi:10.1073/pnas.1711233114 (See also Chapter 26.)

[REF:blacher2019] Blacher E, Bashiardes S, Shapiro H, Rothschild D, Mor U, Dori-Bachash M, et al. (2019). Potential roles of gut microbiome and metabolites in modulating ALS in mice. *Nature*

572(7770):474–480. doi:10.1038/s41586-019-1443-5

[REF:bruggeman2024] Bruggeman A, Vandendriessche C, Hamerlinck H, De Looze D, Tate DJ, Vandeputte D, et al. (2024). Safety and efficacy of faecal microbiota transplantation in patients with mild to moderate Parkinson’s disease (GUT-PARFECT): a double-blind, placebo-controlled, randomised, phase 2 trial. *eClinicalMedicine* 71:102563. doi:10.1016/j.eclinm.2024.102563

[REF:buffington2016] Buffington SA, Di Prisco GV, Auchtung TA, Ajami NJ, Petrosino JF, Costa-Mattioli M (2016). Microbial reconstitution reverses maternal diet-induced social and synaptic deficits in offspring. *Cell* 165(7):1762–1775. doi:10.1016/j.cell.2016.06.001 (Covered in detail in Chapter 14.)

[REF:cryan2022] Cryan JF, Mazmanian SK (2022). Microbiota-brain axis: context and causality. *Science* 376(6597):938–939. doi:10.1126/science.abo4442

[REF:depressionfmtmeta2025] Collaborative authors (2025). Efficacy of faecal microbiota transplantation in major depressive disorder: a systematic review and meta-analysis of 12 randomised controlled trials. *Frontiers in Psychiatry*. (Specific citation to be confirmed during final proofing.)

[REF:depressionfmtmeta2026] Collaborative authors (2026). Faecal microbiota transplantation for depression: a conflicting meta-analysis of seven trials. (Specific citation to be confirmed during final proofing.)

[REF:dominy2019] Dominy SS, Lynch C, Ermini F, Benedyk M, Marczyk A, Konradi A, et al. (2019). *Porphyromonas gingivalis* in Alzheimer’s disease brains: evidence for disease causation and treatment with small-molecule inhibitors. *Science Advances* 5(1):eaau3333. doi:10.1126/sciadv.aau3333 (Cross-referenced with Chapter 11.)

[REF:ferreiro2023] Ferreira AL, Choi J, Ryou J, Newcomer EP, Thompson R, Bollinger RM, et al. (2023). Gut microbiome composition may be an indicator of preclinical Alzheimer’s disease. *Science Translational Medicine* 15(700):eabo2984. doi:10.1126/scitranslmed.abo2984

[REF:gain2021] Detke MJ, et al. (2021). The GAIN trial (atuzaginstat in mild-to-moderate Alzheimer’s disease): a randomized, placebo-controlled, double-blind, phase 2/3 study. (Specific citation to be confirmed during final proofing.)

[REF:heintzbuschart2018] Heintz-Buschart A, Pandey U, Wicke T, Sixel-Döring F, Janzen A, Sittig-Wiegand E, et al. (2018). The nasal and gut microbiome in Parkinson’s disease and idiopathic rapid eye movement sleep behavior disorder. *Movement Disorders* 33(1):88–98. doi:10.1002/mds.27105

[REF:horsager2020] Horsager J, Andersen KB, Knudsen K, Skjærbæk C, Fedorova TD, Okkels N, et al. (2020). Brain-first versus body-first Parkinson’s disease: a multimodal imaging case-control study. *Brain* 143(10):3077–3088. doi:10.1093/brain/awaa238

[REF:hsiao2013] Hsiao EY, McBride SW, Hsien S, Sharon G, Hyde ER, McCue T, et al. (2013). Microbiota modulate behavioral and physiological abnormalities associated with neurodevelopmental disorders. *Cell* 155(7):1451–1463. doi:10.1016/j.cell.2013.11.024

[REF:huang2025] Huang TT, et al. (2025). Psychobiotics in depression: a systematic review of 45 randomized controlled trials. *Nutrients* 17(6). (Specific citation to be confirmed during final proofing.)

- [REF:jangi2016] Jangi S, Gandhi R, Cox LM, Li N, von Glehn F, Yan R, et al. (2016). Alterations of the human gut microbiome in multiple sclerosis. *Nature Communications* 7:12015. doi:10.1038/ncomms12015 (Covered in Chapter 26.)
- [REF:jans2025] Jans M, et al. (2025). A guide to germ-free and gnotobiotic mouse technology to study health and disease. *FEBS Journal*. doi:10.1111/febs.17124
- [REF:jimenezgarcia2024] Jimenez-García AM, et al. (2024). Gut microbiota and cognitive function in Alzheimer's disease: a systematic review and meta-analysis. *Alzheimer's & Dementia: Diagnosis, Assessment & Disease Monitoring*. doi:10.1002/dad2.70057
- [REF:kang2017] Kang DW, Adams JB, Gregory AC, Borody T, Chittick L, Fasano A, et al. (2017). Microbiota transfer therapy alters gut ecosystem and improves gastrointestinal and autism symptoms: an open-label study. *Microbiome* 5(1):10. doi:10.1186/s40168-016-0225-7
- [REF:kang2019] Kang DW, Adams JB, Coleman DM, Pollard EL, Maldonado J, McDonough-Means S, et al. (2019). Long-term benefit of microbiota transfer therapy on autism symptoms and gut microbiota. *Scientific Reports* 9(1):5821. doi:10.1038/s41598-019-42183-0
- [REF:kelly2016] Kelly JR, Borre Y, O'Brien C, Patterson E, El Aidy S, Deane J, et al. (2016). Transferring the blues: depression-associated gut microbiota induces neurobehavioural changes in the rat. *Journal of Psychiatric Research* 82:109–118. doi:10.1016/j.jpsychires.2016.07.019 (Covered in Chapter 13.)
- [REF:killinger2018] Killinger BA, Madaj Z, Sikora JW, Rey N, Haas AJ, Vepa Y, et al. (2018). The vermiform appendix impacts the risk of developing Parkinson's disease. *Science Translational Medicine* 10(465):eaar5280. doi:10.1126/scitranslmed.aar5280
- [REF:mainirekdal2019] Maini Rekdal V, Bess EN, Bisanz JE, Turnbaugh PJ, Balskus EP (2019). Discovery and inhibition of an interspecies gut bacterial pathway for Levodopa metabolism. *Science* 364(6445):eaau6323. doi:10.1126/science.aau6323 (Covered in detail in Chapter 22.)
- [REF:metcalferoach2024] Metcalfe-Roach A, Cirstea MS, Yu AC, Ramay HR, Coker O, Su H, et al. (2024). Metagenomic analysis reveals large-scale disruptions of the gut microbiome in Parkinson's disease. *Movement Disorders* 39(10):1739–1751. doi:10.1002/mds.29959
- [REF:miller2020] Miller AH, Pariante CM (2020). Trial failures of anti-inflammatory drugs in depression. *The Lancet Psychiatry*. (Representative citation for the inflammatory-subtype-of-depression hypothesis; specific citation to be confirmed during final proofing.)
- [REF:needham2022] Needham BD, Funabashi M, Adame MD, Wang Z, Boktor JC, Haney J, et al. (2022). A gut-derived metabolite alters brain activity and anxiety behaviour in mice. *Nature* 602(7898):647–653. doi:10.1038/s41586-022-04396-8
- [REF:postuma2019] Postuma RB, Iranzo A, Hu M, Högl B, Boeve BF, Manni R, et al. (2019). Risk and predictors of dementia and parkinsonism in idiopathic REM sleep behaviour disorder: a multicentre study. *Brain* 142(3):744–759. doi:10.1093/brain/awz030
- [REF:romano2025] Collaborative authors (2025). Machine-learning meta-analysis of Parkinson's disease gut microbiome data: cohort-specific signatures and the generalisability problem. (Specific citation to be confirmed during final proofing.)
- [REF:rothhammer2016] Rothhammer V, Maccanfroni ID, Bunse L, Takenaka MC, Kenison JE, Mayo L, et al. (2016). Type I interferons and microbial metabolites of tryptophan modulate astro-

cyte activity and central nervous system inflammation via the aryl hydrocarbon receptor. *Nature Medicine* 22(6):586–597. doi:10.1038/nm.4106 (Covered in Chapter 26.)

[REF:sampson2016] Sampson TR, Debelius JW, Thron T, Janssen S, Shastri GG, Ilhan ZE, et al. (2016). Gut microbiota regulate motor deficits and neuroinflammation in a model of Parkinson's disease. *Cell* 167(6):1469–1480.e12. doi:10.1016/j.cell.2016.11.018

[REF:scheperjans2015] Scheperjans F, Aho V, Pereira PA, Koskinen K, Paulin L, Pekkonen E, et al. (2015). Gut microbiota are related to Parkinson's disease and clinical phenotype. *Movement Disorders* 30(3):350–358. doi:10.1002/mds.26069

[REF:sgritta2018] Sgritta M, Dooling SW, Buffington SA, Momin EN, Francis MB, Britton RA, Costa-Mattioli M (2018). Mechanisms underlying microbial-mediated changes in social behavior in mouse models of autism spectrum disorder. *Neuron* 101(2):246–259.e6. doi:10.1016/j.neuron.2018.11.018 (Covered in detail in Chapter 14.)

[REF:sharon2019] Sharon G, Cruz NJ, Kang DW, Gandal MJ, Wang B, Kim YM, et al. (2019). Human gut microbiota from autism spectrum disorder promote behavioral symptoms in mice. *Cell* 177(6):1600–1618.e17. doi:10.1016/j.cell.2019.05.004

[REF:tzeng2018] Tzeng NS, Chung CH, Lin FH, Chiang CP, Yeh CB, Huang SY, et al. (2018). Anti-herpetic medications and reduced risk of dementia in patients with herpes simplex virus infections — a nationwide, population-based cohort study in Taiwan. *Neurotherapeutics* 15(2):417–429. doi:10.1007/s13311-018-0611-x

[REF:vallescolomer2019] Valles-Colomer M, Falony G, Darzi Y, Tigchelaar EF, Wang J, Tito RY, et al. (2019). The neuroactive potential of the human gut microbiota in quality of life and depression. *Nature Microbiology* 4(4):623–632. doi:10.1038/s41564-018-0337-x (Covered in Chapter 13.)

[REF:vogt2017] Vogt NM, Kerby RL, Dill-McFarland KA, Harding SJ, Merluzzi AP, Johnson SC, et al. (2017). Gut microbiome alterations in Alzheimer's disease. *Scientific Reports* 7(1):13537. doi:10.1038/s41598-017-13601-y

[REF:wallen2022] Wallen ZD, Demirkan A, Twa G, Cohen G, Dean MN, Standaert DG, et al. (2022). Metagenomics of Parkinson's disease implicates the gut microbiome in multiple disease mechanisms. *Nature Communications* 13(1):6958. doi:10.1038/s41467-022-34667-x

[REF:xue2023] Xue LJ, Yang XZ, Tong Q, Shen P, Ma SJ, Wu SN, et al. (2023). Fecal microbiota transplantation in patients with Parkinson's disease: clinical trial results from a randomized, placebo-controlled design. *Gut Microbes* 15(2):2284247. doi:10.1080/19490976.2023.2284247

[REF:xu1989] Xu Y, Sladky JT, Brown MJ (1989). Dose-dependent expression of neuronopathy after experimental pyridoxine intoxication. *Neurology* 39(8):1077–1083. doi:10.1212/wnl.39.8.1077

[REF:hoover1981] Hoover DM, Carlton WW (1981). The subacute neurotoxicity of excess pyridoxine HCl and clioquinol in beagle dogs. II. Pathology. *Veterinary Pathology* 18(6):757–768. doi:10.1177/030098588101800606

[REF:alvarez2026] Alvarez M, Poveda S, Cisneros A, Parra D, Luna M, Rincón O, Guzman I (2026). B vitamin deficiencies and associated neuropathies. *Current Nutrition Reports* 15(1):8. doi:10.1007/s13668-025-00723-3

[REF:snodgrass1992] Snodgrass SR (1992). Vitamin neurotoxicity. *Molecular Neurobiology* 6(1):41–73. doi:10.1007/BF02935566

- [REF:vrolijk2017] Vrolijk MF, Opperhuizen A, Jansen EHJM, Hageman GJ, Bast A, Haenen GRMM (2017). The vitamin B6 paradox: supplementation with high concentrations of pyridoxine leads to decreased vitamin B6 function. *Toxicology In Vitro* 44:206–212. doi:10.1016/j.tiv.2017.07.009
- [REF:hadtstein2021] Hadtstein F, Vrolijk M (2021). Vitamin B-6-induced neuropathy: exploring the mechanisms of pyridoxine toxicity. *Advances in Nutrition* 12(5):1911–1929. doi:10.1093/advances/nmab033
- [REF:reddy2021] Reddy P (2021). Preventing vitamin B6-related neurotoxicity. *American Journal of Therapeutics* 29(6):e637–e643. doi:10.1097/MJT.0000000000001460
- [REF:magnusdottir2015] Magnúsdóttir S, Ravcheev D, de Crécy-Lagard V, Thiele I (2015). Systematic genome assessment of B-vitamin biosynthesis suggests co-operation among gut microbes. *Frontiers in Genetics* 6:148. doi:10.3389/fgene.2015.00148
- [REF:radjabzadeh2020] Radjabzadeh D, Boer CG, Beth SA, van der Wal P, Kiefte-De Jong JC, Jansen MAE, et al. (2020). Diversity, compositional and functional differences between gut microbiota of children and adults. *Scientific Reports* 10(1):1040. doi:10.1038/s41598-020-57734-z
- [REF:wu2025] Wu Y, Ling Z, Peng L, Wang X, He Y, Khan A, et al. (2025). Vitamin B6 form produced by *Lactobacillus* induces metabolic disorder and suppresses multi-pathogenic bacteria. *Communications Biology* 9(1):82. doi:10.1038/s42003-025-09348-2
- [REF:abdollahiyan2023] Abdollahiyan S, Nabavi-Rad A, Keshavarz Azizi Raftar S, Monnoye M, Salarieh N, Farahanie A, et al. (2023). Characterization of gut microbiome composition in Iranian patients with nonalcoholic fatty liver disease and nonalcoholic steatohepatitis. *Scientific Reports* 13(1):20584. doi:10.1038/s41598-023-47905-z
- [REF:gut_dysbiosis_b6_2024] (2024). Gut dysbiosis-induced vitamin B6 metabolic disorder contributes to chronic stress-related abnormal behaviors in a cortisol-independent manner. *Gut Microbes* 17(1):2447824. doi:10.1080/19490976.2024.2447824

Chapter 30: The Microbiome as a Diagnostic Tool

30.0 Reading the Ecosystem

Imagine a clinic visit ten years from now. You arrive for a routine check-up. Alongside the familiar blood draw, the nurse hands you a small tube and asks you to provide a stool sample — or perhaps you have already provided one, passively, via a sensor mounted in your toilet at home. A few days later, the results arrive. Your blood panel is normal. Your microbiome report is not. It flags a shift in your gut community — a cluster of oral-origin bacteria that have established themselves in your colon, a depletion of the butyrate producers your doctor has been monitoring, and an algorithm-generated risk score suggesting a threefold elevation in your likelihood of developing colorectal cancer within the next five years. Your doctor orders a colonoscopy. A small adenoma is found and removed. Cancer prevented.

Is this science fiction? In 2026, largely yes. No microbiome-based diagnostic test has been approved by the US Food and Drug Administration. No clinical guideline in any country recommends routine microbiome testing for disease screening. An international consensus statement published in the *Lancet Gastroenterology & Hepatology* in late 2024, authored by an expert panel spanning gastroenterology, microbiology, bioinformatics, and clinical medicine, concluded that “the time is not yet ripe to use microbiota analysis widely in clinical practice, particularly for the early diagnosis of many diseases or as a guide to treatment” [REF:icc2024]. That is the honest starting point for this chapter.

And yet. The research literature now contains dozens of studies — many of them large, multi-cohort, and independently validated — reporting that microbiome-based classifiers can detect specific diseases with an accuracy that matches or exceeds existing clinical tests. In colorectal cancer, stool microbiome signatures achieve diagnostic accuracies comparable to the faecal immunochemical test that is the backbone of current screening programmes. In inflammatory bowel disease, a microbiome-based diagnostic has outperformed faecal calprotectin, the most widely used non-invasive marker. In liver cirrhosis, gut microbial signatures discriminate between healthy and cirrhotic patients with an accuracy above ninety per cent. And in Parkinson’s disease, microbiome changes are now detectable in people who carry genetic risk factors but have not yet developed any symptoms — raising the extraordinary possibility that a stool test could flag neurodegeneration years before the first tremor.

The gap between what the research shows and what the clinic uses is one of the most instructive gaps in contemporary medicine. Understanding why it exists — and what it will take to close it — is the purpose of this chapter. The answer involves technical barriers (can we make the tests cheap, fast, and reproducible?), scientific barriers (do the results hold up when you test them in new populations?), regulatory barriers (when is a microbiome test a medical device?), and commercial distortions (what happens when companies sell microbiome “insights” before the science supports them?). Each section of this chapter takes up a different piece of this puzzle, moving from the diagnostic applications closest to clinical use through intermediate-term prospects to the further-future technologies that may, within the reader’s lifetime, make the opening scenario of this section look conservative.

30.1 Microbial Biomarkers for Disease Screening

The logic of microbial diagnostics

Before we look at specific diseases, it is worth pausing to ask why the microbiome should be diagnostic at all. The answer lies in a principle that has been building throughout this book: the microbiome is not a passive bystander in human disease. It is an active participant, shaped by the disease process and, in many cases, shaping it in return. When a tumour grows in the colon, the bacteria that thrive in its immediate vicinity change — some because the tumour provides them with new nutrients, others because the tumour’s disruption of the mucosal barrier creates new ecological niches. When the immune system is chronically inflamed, as in inflammatory bowel disease, the microbial community shifts toward species that tolerate or exploit inflammation. When the liver fails, the gut barrier weakens, and bacteria that normally live in the mouth — species that have no business in the colon — begin to appear in faecal samples, a phenomenon microbiologists have taken to calling the “oralisation” of the gut.

These shifts are not subtle. They are often large enough to be detected by sequencing a stool sample and feeding the data into a machine learning algorithm. The question is not whether the shifts exist — they do, reproducibly, across dozens of studies. The question is whether detecting them adds anything useful to the diagnostic tools we already have.

That question has different answers for different diseases. In what follows, we take the cases one by one, starting with the one where the answer is closest to “yes.”

Colorectal cancer: the flagship

Colorectal cancer is the most advanced diagnostic use-case in microbiome science. There are good reasons for this. CRC is common — the third most frequently diagnosed cancer worldwide, killing nearly a million people each year. It develops slowly, usually over a decade or more, through a well-characterised progression from normal mucosa to adenomatous polyp to carcinoma. That slow progression makes it an ideal target for screening: catch the polyp, remove the polyp, prevent the cancer. Existing screening tools — colonoscopy, faecal immunochemical testing (FIT), and the newer multi-target stool DNA tests — save lives but are imperfect. Colonoscopy is invasive, expensive, and requires bowel preparation that deters many patients from completing it. FIT is cheap and non-invasive but misses a substantial fraction of advanced adenomas and some early-stage cancers. There is room for improvement.

The microbiome entered this story in 2014, when Georg Zeller and colleagues at the European Molecular Biology Laboratory published the first study showing that faecal microbiome data could detect CRC with reasonable accuracy [REF:zeller2014]. But the field truly matured in 2019, when two landmark meta-analyses — published simultaneously in *Nature Medicine* by competing groups who had, by remarkable scientific coincidence, been working on the same question with overlapping datasets — established that the signal was not an artefact of any single population, sequencing protocol, or analytical pipeline.

The first, led by Jakob Wirbel in Peer Bork’s group at EMBL, pooled eight geographically diverse faecal shotgun metagenomic studies encompassing 768 samples [REF:wirbel2019]. They identified a core set of 29 bacterial species significantly enriched in CRC metagenomes. The enriched

species were not random. Many were oral-origin organisms — bacteria that normally reside in the mouth but had somehow colonised the colorectal tumour environment. *Fusobacterium nucleatum*, the species we met in Chapter 28 as a direct promoter of colorectal carcinogenesis, was the most prominent. But it was not alone: *Porphyromonas asaccharolytica*, *Parvimonas micra*, *Peptostreptococcus stomatis*, and *Solobacterium moorei* all showed consistent enrichment across cohorts on three continents and in seven countries. When the team trained machine learning classifiers on data from one cohort and tested them on another — the gold standard for assessing whether a diagnostic signature is robust — the classifiers maintained their accuracy. Training on multiple cohorts simultaneously improved performance further.

The second meta-analysis, led by Andrew Maltez Thomas in Nicola Segata's group at the University of Trento, included five publicly available datasets plus two new cohorts — 969 faecal metagenomes in total [REF:thomas2019]. Their predictive signatures achieved an average area under the receiver operating characteristic curve (AUC) of 0.84 — a number worth pausing over. An AUC of 0.50 is a coin flip; 1.00 is a perfect test. The standard faecal immunochemical test, FIT, achieves an AUC of roughly 0.80 for detecting CRC. So the microbiome-based classifier, using nothing but the bacterial DNA in a stool sample, was already performing in the same range as the screening test currently used by millions of people worldwide. Thomas's group also uncovered a novel metabolic link: the gene encoding choline trimethylamine-lyase (cutC) was significantly over-abundant in CRC metagenomes. This enzyme converts dietary choline into trimethylamine — a precursor to TMAO, the cardiovascular risk molecule we met in Chapter 27. The finding suggested that the same microbial metabolic pathway might be relevant to both cancer and heart disease, a cross-disease connection that had not been anticipated.

These twin papers were a turning point. They established three things that earlier, smaller studies could not. First, the CRC microbiome signature is globally robust — it survives differences in diet, geography, ethnicity, and laboratory protocol. Second, the signal is multi-species rather than single-organism: *Fusobacterium nucleatum* is the poster child, but detecting CRC from the microbiome requires a community-level analysis. Third, the diagnostic performance is clinically meaningful — not yet superior to the best existing tests, but comparable, and potentially complementary.

The field has not stood still since 2019. In 2025, a massive pooled analysis expanded the evidence base to 3,741 stool metagenomes from 18 cohorts [REF:pooled2025]. This study — the largest of its kind — did something the earlier meta-analyses could not: it moved from species-level to strain-level resolution. Different clades of *Fusobacterium nucleatum* were shown to have different associations with CRC — a finding that mirrors the cancer biology we discussed in Chapter 28, where Susan Bullman's group demonstrated that a specific clade of *F. nucleatum* (Fna C2) dominates the colorectal cancer niche [REF:bullman2024]. The pooled analysis also identified cross-stage biomarkers: species whose abundance changes not just between healthy controls and cancer patients, but between normal mucosa, adenoma (precancerous polyp), and carcinoma. This is diagnostically important because the greatest benefit of screening lies in catching the adenoma — the polyp that has not yet become cancer — and the 2025 study suggests that microbial signatures can, to some degree, make that distinction.

Meanwhile, in Geneva, a team led by Mirko Tričković produced the first detailed catalogue of gut bacteria at the subspecies level, applying it to CRC detection [REF:trickovic2025]. Using AI-based analysis of subspecies diversity, they identified subtle microbial patterns linked to cancer that species-level analysis had missed. The analogy is instructive: imagine trying to solve a crime from a photograph taken at a distance. You can see that someone was in the room. Subspecies-level analysis is like zooming in until you can read the number on their badge.

So the evidence is strong. Why is there no microbiome-based CRC screening test in clinical use?

The answer has several parts. First, any new screening test must be compared not just against FIT but against the best available alternative. In 2024, Exact Sciences published data on Cologuard Plus, a next-generation multi-target stool DNA test that combines human DNA methylation markers with FIT, achieving a sensitivity of ninety-four per cent for CRC and a specificity for advanced neoplasia above ninety per cent in a study of more than 20,000 participants. That is a very high bar. The microbiome-based approach, with an AUC of 0.85, is competitive but not superior. The obvious next step — combining microbial markers with human DNA markers and FIT in a single test — has not yet been attempted in a large prospective trial.

Second, there is the reproducibility problem. The cross-cohort classifiers in the Wirbel and Thomas papers held up impressively, but they were validated on research cohorts that had been carefully collected, stored, and processed. A clinical screening test must work on samples collected by patients at home, mailed through the post (sometimes in summer heat), processed in commercial laboratories with varying protocols, and interpreted by clinicians with no bioinformatics training. The gap between research-grade and clinical-grade performance has killed many a promising biomarker.

Third, there is the regulatory pathway. A stool-based microbiome test for CRC screening would likely need to be approved as an in-vitro diagnostic device — a process that requires not just evidence of analytical validity (does the test measure what it claims to measure?) and clinical validity (does the result correlate with the disease?), but also clinical utility (does using the test improve patient outcomes compared with not using it?). Demonstrating clinical utility requires a prospective trial in which patients are randomised to microbiome-based screening versus standard screening, and followed for years to see whether the new test catches cancers earlier, reduces cancer deaths, or reduces unnecessary colonoscopies. No such trial has been conducted.

The honest summary: the microbiome can detect colorectal cancer in a stool sample with an accuracy that is clinically meaningful and globally robust. Whether it should be used for screening — and whether it adds anything to the tests we already have — remains to be demonstrated in the kind of large, prospective, outcome-driven trial that regulators and clinicians rightly demand. The trajectory is clear. The destination has not yet been reached.

Inflammatory bowel disease: outperforming the standard

If CRC is the flagship application of microbiome diagnostics, inflammatory bowel disease is where the field has taken the most decisive step toward clinical translation.

The current non-invasive diagnostic workhorse for IBD is faecal calprotectin — a protein released by neutrophils when they migrate into the intestinal lumen during inflammation. Calprotectin is useful: it is cheap, widely available, and reliably elevated in active IBD. But it has significant limitations. It is not specific to IBD — it rises in infections, during NSAID use, and in other inflammatory conditions. It cannot distinguish ulcerative colitis from Crohn's disease, a distinction that matters because the two conditions have different treatment pathways, different surgical options, and different long-term prognoses. And it cannot tell a clinician whether a patient who is currently in remission is likely to relapse. Definitive diagnosis still requires colonoscopy with biopsy.

In 2024, a study from the Chinese University of Hong Kong, led by the Microbiota I-Center and published in *Nature Medicine*, changed the landscape [REF:zheng2024]. Jiaying Zheng, Siew Ng, and colleagues assembled an extraordinary dataset: 5,979 faecal metagenomes from 13 cohorts across eight countries and eight ethnicities. The diversity of this dataset was deliberate. A diagnostic test

that works only in Hong Kong Chinese patients is not clinically useful. A diagnostic test that works across populations, diets, and genetic backgrounds is.

From their discovery cohort of 609 samples, the team identified ten bacterial species whose combined abundances could distinguish ulcerative colitis from healthy controls, and nine species that could distinguish Crohn's disease from healthy controls. Random forest classifiers built on these species achieved an AUC of 0.90 in the discovery cohort — an excellent result, but one that might be suspected of overfitting. The critical test was what happened when they applied the classifier to the remaining cohorts: across all validation populations, the AUC remained at 0.82 for IBD as a whole. For ulcerative colitis specifically, the classifier performed at 0.88 in discovery and 0.89 in the broadest validation set. For Crohn's disease, it achieved 0.87 in discovery, though cross-ethnic validation was more variable (0.73–0.75 in some populations). These are not perfect numbers, but they are clinically competitive.

The study's most important contribution, however, was not the metagenomics classifier. It was the translation step. Recognising that shotgun metagenomic sequencing is too expensive and too slow for routine clinical use, the team developed a multiplex droplet digital PCR (ddPCR) assay targeting the same diagnostic species. Droplet digital PCR is a well-established clinical technology — fast, relatively cheap, and available in diagnostic laboratories worldwide. Their ddPCR-based diagnostic model achieved an AUC of 0.85 for ulcerative colitis and 0.87 for Crohn's disease — and in head-to-head comparison, it numerically outperformed faecal calprotectin.

This matters because it demonstrates a viable pathway from discovery to clinic. The typical trajectory of a microbiome biomarker study is: discover the signal using expensive metagenomics, publish the paper, and then watch the finding languish because no clinical laboratory is equipped to perform shotgun sequencing on every stool sample. The Zheng team short-circuited this trajectory by building, within the same study, a PCR-based assay that any diagnostic lab could run. If their results are validated in a prospective clinical trial — and that trial has not yet been conducted — this could become the first microbiome-based test to enter routine gastroenterological practice.

Liver disease: the gut as a window onto a distant organ

The liver has no microbiome of its own — at least not in health. Unlike the gut, the skin, or the mouth, the liver in a healthy person is essentially sterile, protected by the portal venous system and the hepatic immune surveillance that filters the blood arriving from the intestine. When microbiologists talk about “the liver microbiome,” they are usually talking about what the gut microbiome can tell us about the liver.

That turns out to be a great deal. As we discussed in Chapter 15, the gut and the liver are connected by the portal vein — a direct vascular highway that carries the products of intestinal digestion, microbial metabolism, and, when the gut barrier weakens, the bacteria themselves, straight to the liver. In chronic liver disease, this connection becomes pathological. The gut barrier degrades. Bacterial products — lipopolysaccharide, peptidoglycan, flagellin — translocate to the liver in increasing quantities, driving inflammation and fibrosis. And the gut microbial community itself changes in characteristic ways, shifting toward species that tolerate the altered bile acid profile and increased intestinal permeability of the cirrhotic gut.

These shifts are large enough to be diagnostic. In 2019, Cyrielle Caussy and colleagues at the University of California San Diego built a random forest classifier using faecal microbiome data from patients with NAFLD-associated cirrhosis and demonstrated an AUC of 0.92 for detecting cirrhosis

in their derivation cohort [REF:caussy2019]. A panel of thirty features — twenty-seven of them bacterial taxa — was sufficient. In 2020, a separate study combined metagenomic and metabolomic signatures to achieve an AUC of 0.91 for cirrhosis detection, and showed that combining the metagenomic signature with just two clinical variables — age and serum albumin — accurately distinguished cirrhotic from non-cirrhotic patients even across geographically separated populations [REF:iebbba2020].

What makes the liver disease story distinctive is the ecological pattern it reveals. In cirrhosis, the gut microbiome undergoes what researchers have called “oralisation”: species that normally reside in the mouth — *Veillonella*, *Streptococcus*, certain *Lactobacillus* species — begin to appear in the colon in substantial numbers, while typical colonic genera like *Bacteroides* and *Prevotella* decline. This is thought to reflect the failure of gastric acid and bile acid barriers that normally prevent oral bacteria from surviving the journey to the lower gut. In a sense, the cirrhotic gut microbiome is a map of the body’s failing defences — and reading that map from a stool sample is cheaper, faster, and less invasive than a liver biopsy.

Parkinson’s disease: screening before the tremor

In Chapter 29, we traced the Parkinson’s disease–microbiome story from the Braak hypothesis through the prodromal dysbiosis studies to the GUT-PARFECT trial. We noted that gut microbiome signatures can be detected in patients with REM sleep behaviour disorder — the most specific prodromal marker of Parkinson’s — years before the motor symptoms that define the clinical diagnosis. That observation, interesting in the context of disease biology, becomes profound in the context of diagnostics.

In April 2026, a study from Elisa Menozzi and Anthony Schapira at UCL Queen Square Institute of Neurology, published in *Nature Medicine* as part of the ASAP Collaborative Network, brought this possibility into sharp focus [REF:menozzi2026]. The team studied 271 patients with established Parkinson’s disease, 150 healthy controls, and — crucially — 43 individuals who carry pathogenic variants in the GBA1 gene. GBA1 variants are the most common genetic risk factor for Parkinson’s disease: carriers have a substantially elevated lifetime risk of developing the disease, but many never will. They are, in effect, a population in whom the disease may or may not be silently developing.

The results were striking. Approximately twenty-five per cent of the gut microbiome in the GBA1 carriers — who had no symptoms of Parkinson’s disease — was intermediate between the healthy controls and the Parkinson’s patients. The microbiome of these pre-symptomatic carriers was already shifting toward the Parkinson’s pattern, even though they showed no tremor, no rigidity, no clinical signs at all. Moreover, the degree of this intermediate shift correlated with prodromal symptoms — subtle autonomic and sensory changes that, with hindsight, often presage Parkinson’s by years. The finding was validated in three independent cohorts from the United States, Korea, and Turkey, spanning 638 Parkinson’s patients and 319 healthy controls.

The implication is extraordinary: a stool test could, in principle, identify people at elevated risk of Parkinson’s disease years or even decades before the diagnosis. But this implication immediately collides with a difficult ethical question that links forward to Chapter 31. As of 2026, there is no proven intervention that can prevent Parkinson’s disease in an at-risk individual. We cannot slow the progression. We cannot stop the α -synuclein from misfolding. If a stool test tells a fifty-year-old woman that her gut microbiome looks like an early-Parkinson’s microbiome, what do we do with that information? The GUT-PARFECT trial (Chapter 29) showed that FMT can produce a

modest improvement in motor symptoms in established Parkinson's disease, but no trial has tested whether microbiome intervention in the prodromal phase can prevent the disease from developing. The Menozzi study opens a door. What lies beyond it is not yet clear.

The cross-cohort generalisability problem is also present here. A 2025 machine learning meta-analysis of nearly 4,500 Parkinson's disease microbiome samples found that within-study classifiers achieved a mean AUC of about seventy-two per cent — decent but not outstanding [REF:romano2025]. When the same classifiers were tested across cohorts, performance fell to sixty-one per cent — barely better than a coin flip. The signal is real but tangled with the same confounders — diet, geography, ethnicity, sequencing protocol — that bedevil every microbiome diagnostic application. Whether the microbial-fingerprint normalisation approach that Jiyoung Ahn's group introduced for melanoma immunotherapy prediction (Chapter 28) can rescue portability for Parkinson's is an open question.

The oral microbiome as a diagnostic window

In Chapter 9, we introduced the idea that the mouth might serve as a diagnostic window onto systemic disease. The oral microbiome is the easiest human microbiome to sample — a saliva collection requires no bowel preparation, no stool handling, and no patient reluctance. If oral bacteria carry diagnostic information about diseases beyond the mouth, saliva-based diagnostics could be transformative.

The evidence is accumulating, though it remains largely at the association stage. The most tantalising connection is with pancreatic cancer — one of the deadliest and hardest-to-detect malignancies, with a five-year survival rate below ten per cent. A 2025 systematic review and meta-analysis assessed the diagnostic accuracy of the oral microbiome for pancreatic cancer and found consistent associations with *Porphyromonas gingivalis* and *Fusobacterium nucleatum* — the same species that feature in the colorectal cancer story [REF:oralpc2025]. Additional associations have been reported with *Alloprevotella* and several other taxa. The biological plausibility is not unreasonable: *P. gingivalis* has been shown to activate pathways that promote pancreatic cell proliferation, and oral bacteria can reach the pancreas via the biliary tree or the bloodstream. But no saliva-based pancreatic cancer test has progressed beyond the association stage to clinical validation.

A different and potentially broader application emerged in a 2026 *Nature Communications* study showing that oral microbiome data can predict biological age — and, more importantly, that an “Oral Microbiome Aging Acceleration Score” independently predicted all-cause mortality and frailty, and enhanced cancer risk prediction, in a large cohort [REF:oralage2026]. The researchers identified sixty-four bacterial genera whose abundances change systematically with age and built a machine learning model that predicted chronological age from oral microbiome data alone. People whose oral microbiome made them look older than their calendar age — those with high “aging acceleration” — had worse health outcomes. Whether this score has clinical utility beyond what can be achieved with simpler aging biomarkers (grip strength, walking speed, standard blood markers) remains to be seen. But the principle is provocative: your mouth bacteria know how old you really are.

Multi-cancer blood-based screening: a cautionary tale

Every diagnostic success story in microbiome science has a counterpart in failure, and this chapter would be dishonest if it told only the successes. The most spectacular failure in the field — and

the one that carries the most important lesson — is the story of circulating microbial DNA as a multi-cancer diagnostic.

The idea is seductive. Tumours, as they grow, shed not only human DNA into the bloodstream (the basis of the rapidly growing “liquid biopsy” field) but also microbial DNA — fragments from the bacteria and viruses that colonise or inhabit the tumour microenvironment. If different cancers harbour different microbial communities (which they do — Chapter 28), then the microbial DNA circulating in the blood might carry a tumour-specific signature. A blood test that could screen for multiple cancers simultaneously, using microbial rather than human DNA, would be a diagnostic revolution.

In March 2020, Gregory Poore, Rob Knight, and colleagues at the University of California San Diego published exactly this claim in *Nature* [REF:poore2020_RETRACTED]. They re-analysed whole-genome and whole-transcriptome sequencing data from The Cancer Genome Atlas — 18,116 samples spanning thirty-three cancer types — and reported unique microbial signatures in tissue and blood that could distinguish between most major cancers. Machine learning models trained on blood microbial DNA could identify a patient’s cancer type. The paper was widely covered in the press. Micronoma, a company co-founded by Knight and two other authors, raised millions of dollars to develop a blood test for lung cancer based on the approach.

Then the finding collapsed. In 2023, Abraham Gihawi and colleagues published a devastating re-analysis showing that the microbial signatures reported in the paper were largely artefacts [REF:gihawi2023]. The microbial databases used to classify the sequencing reads were contaminated with mislabelled human DNA — sequences that had snuck into microbial reference genomes during database construction. When a human tumour sample was screened against these contaminated databases, human DNA fragments were misidentified as microbial, producing what appeared to be cancer-specific microbial signatures but were in fact technical noise. The signatures that appeared to distinguish cancer types were driven by batch effects and database contamination, not by genuine biology.

In June 2024, *Nature* retracted the paper. A subsequent *Nature Communications* paper that had built on the same approach — predicting cancer prognosis and drug response from the tumour microbiome — was also retracted in 2025. Micronoma, the company that had sought to commercialise the approach, crumbled.

This is not a story about scientific fraud. It is a story about the devastating consequences of inadequate bioinformatics controls when working with low-biomass microbial signals in human tissues. The microbial DNA that circulates in blood is present at vanishingly low concentrations relative to human DNA — a needle in a haystack of human sequences. Separating genuine microbial signal from contamination, misclassification, and batch effects is one of the hardest technical problems in the field. Poore’s group was not the first to be tripped up by it, and they will not be the last.

An analogy may help. Imagine trying to hear a single flute in an orchestra of a thousand instruments, through a recording that has background static. The flute — the microbial DNA — is real. But the static — contamination, database errors, misclassified sequences — can easily be mistaken for the flute. The Poore paper heard static and called it music. That does not mean the flute is not there. Other groups, working with more rigorous contamination controls, have reported promising results. A 2024 study in *Cell Reports Medicine* used circulating microbial DNA to detect lung cancer with a sensitivity of nearly eighty-eight per cent and an AUC of 0.93 in an independent validation cohort, including detection of tumours smaller than one centimetre [REF:liang2024]. But the bar for believing these claims is now, rightly, much higher than it was before the retraction.

The Poore episode carries three lessons for the reader. First: extraordinary diagnostic claims require extraordinary bioinformatics rigour, and “published in *Nature*” is not, by itself, a guarantee of either. Second: the pathway from a promising research finding to a validated clinical test is long, and companies that commercialise findings before that pathway is complete are building on sand. Third: the self-correcting mechanism of science — re-analysis, replication attempts, post-publication critique — works, but it works slowly. Three years elapsed between the original publication and the retraction. During that time, the finding was cited hundreds of times and influenced both research directions and investment decisions. The reader who encounters the next “blood test for cancer” headline based on microbial DNA should ask: have the contamination controls been independently verified?

30.2 Companion Diagnostics: Predicting Drug Response

The diagnostic question shifts in this section from “do you have a disease?” to “which treatment will work for you?” This is the domain of companion diagnostics — tests performed before or alongside treatment to predict whether a specific therapy will be effective for a specific patient. In oncology, companion diagnostics are already routine: a breast cancer patient’s tumour is tested for HER2 expression before trastuzumab is prescribed; a lung cancer patient’s tumour is tested for PD-L1 expression before immunotherapy is offered. The question is whether the microbiome can serve a similar function.

Immunotherapy: the strongest case

We built the foundation for this section in Chapter 28, where we traced the evidence that gut microbiome composition influences the efficacy of immune checkpoint inhibitors — the cancer drugs that work by releasing the brakes on the patient’s own immune system. The landmark papers by Gopalakrishnan, Matson, and Routy in 2018 established that patients with certain gut microbial profiles respond to anti-PD-1 therapy while others do not. The Usyk/Ahn microbial-fingerprint study in 2026 showed that matching patients by gut community structure, rather than geography, dramatically improved prediction of melanoma recurrence under immunotherapy [REF:usyk2026]. The logical next step is to turn this research into a clinical test: sequence the patient’s stool before starting treatment, predict whether they will respond, and use that prediction to guide therapy.

The first commercial attempt to do this is BiomeOne, a stool-based microbiome algorithm developed to predict response to first-line immune checkpoint inhibitors across multiple cancer types [REF:biomeone2023]. In a study encompassing non-small-cell lung cancer, renal cell carcinoma, and metastatic melanoma, BiomeOne classified patients into groups by likelihood of clinical benefit before treatment began, using an “immunotherapy-favourable microbiome signature.” The authors reported that BiomeOne outperformed PD-L1 expression — the current companion diagnostic for immunotherapy — in predicting response. If this result holds up, it would mean that a stool test is more informative than a tumour biopsy for deciding whether a patient should receive immunotherapy.

But “if this result holds up” is doing heavy work in that sentence. BiomeOne has been evaluated in a single-company study. It has not been independently validated in a multi-centre trial with blinded outcome assessment. The sample sizes were modest. And the history of companion diagnostics is littered with biomarkers that looked promising in initial studies and failed in prospective validation. The regulatory pathway for a companion diagnostic is, if anything, more demanding than for a

screening test: the test must not only predict response accurately but must be shown to improve patient outcomes when used to guide treatment decisions. That requires a randomised trial in which patients are assigned to microbiome-guided therapy versus standard therapy, and such a trial has not been conducted.

The research is moving in the right direction. A 2025 genome-resolved metagenomic analysis of 624 stool samples from eleven cohorts across five cancer types and three continents identified *Blautia* species — producers of acetate via the Wood-Ljungdahl pathway — as conserved biomarkers of immunotherapy response [REF:blautia2025]. A separate study showed that enriched metabolic pathways in the gut microbiome predicted ICI response across demographic regions and cancer types, suggesting that functional (pathway-level) signatures may be more portable across populations than taxonomic (species-level) ones. This echoes a theme from Chapter 28: asking “what are the bacteria doing?” may be more diagnostically useful than asking “which bacteria are there?”

Beyond immunotherapy: the pharmacomicrobiomics vision

The immunotherapy story is the most advanced, but it is a specific instance of a general principle: the microbiome sits between the patient and virtually every orally administered drug, modifying both efficacy and toxicity in ways that traditional pharmacology has largely ignored. In Chapter 22, we explored several examples. Metformin, the most widely prescribed diabetes drug in the world, exerts at least part of its glucose-lowering effect through the gut microbiome. Levodopa, the cornerstone of Parkinson’s treatment, is degraded by certain strains of *Enterococcus faecalis* before it can reach the brain. Proton pump inhibitors reshape the gut community in ways that may increase susceptibility to *Clostridioides difficile* infection.

The diagnostic vision that emerges from these observations is a pre-treatment stool test that profiles a patient’s drug-metabolising bacterial enzymes and predicts, before the first dose is taken, how efficiently the drug will be absorbed, how much will be inactivated in the gut, and what side effects are most likely. For levodopa, such a test might identify Parkinson’s patients who carry high levels of the tyrosine decarboxylase gene in their gut community — patients who will need higher doses, or co-administration of a bacterial enzyme inhibitor, to achieve therapeutic brain levels. For metformin, such a test might predict which patients will experience the gastrointestinal side effects that cause a significant minority to discontinue the drug.

This is a compelling vision. It is also, as of 2026, almost entirely theoretical. No pharmacomicrobiomics companion diagnostic has been developed for any drug outside immunotherapy. The Maini Rekdal paper on levodopa metabolism (Chapter 22) was published in 2019; seven years later, no clinical test for TDC-carrying bacteria has been commercialised. The reasons are mundane but important: developing a diagnostic requires not just identifying the relevant enzyme or pathway but validating its predictive value in a prospective clinical cohort, demonstrating that the test is analytically reproducible, and navigating the regulatory pathway. Each of these steps takes years and costs millions of dollars. The science is ready. The infrastructure is not.

30.3 Sensors, Capsules, and Real-Time Monitoring

Everything described so far in this chapter relies on the same fundamental workflow: collect a sample (stool, saliva, blood), send it to a laboratory, extract and sequence the DNA, run an algorithm, and return a result. This process takes days to weeks and provides a snapshot — a single

time-point picture of a microbial community that is, as we have emphasised throughout this book, continuously shifting. The technologies described in this section aim to change that — to move from snapshots to movies, from laboratory processing to real-time sensing, and from periodic testing to continuous monitoring.

Smart toilets: what they can and cannot do

The most natural point of contact between the human body and a monitoring device is, arguably, the toilet. We use one every day. We produce a sample that contains, among other things, a concentrated representation of our gut microbial community. It is not surprising that engineers have tried to turn the toilet into a diagnostic instrument.

The first serious academic prototype was published in *Nature Biomedical Engineering* in 2020 by Seung-min Park and colleagues at Stanford [REF:park2020]. Their “mountable toilet system” used pressure and motion sensors, biometric identification (including, memorably, anal-print recognition for user identification — a detail that generated considerable media attention), urinalysis strips, a computer-vision uroflowmeter, and a stool-imaging system. The device could track stool consistency against the Bristol Stool Scale, detect occult blood, and monitor urinary biomarkers. It was, in essence, a continuous health-monitoring platform disguised as a bathroom fixture.

By 2025–2026, commercial products have begun to appear. Kohler Health’s Dekoda, shipping in late 2025 at \$599, uses optical sensors and machine learning to assess hydration status, gut health indicators, and blood presence. Throne One, shipping in early 2026 at \$399, uses computer vision for stool consistency and colour tracking, paired with a meal-tracking app. A 2025 *Nature Protocols* paper described a cloud-based passive defecation monitoring system for continuous gut health surveillance.

These are genuinely useful devices for tracking broad trends in digestive health. But the reader should note what they are not doing: they are not analysing the microbiome. What current smart toilets can measure — stool colour, consistency, frequency, blood, and certain urinary metabolites — are proxies for gut health, not direct readouts of microbial community composition. Accurate microbiome profiling requires cold-chain sample storage, bacterial DNA extraction, sequencing, and computationally intensive bioinformatic analysis. Automating and miniaturising that entire pipeline — from a fresh stool sample to a microbial profile — inside or alongside a toilet is currently beyond our engineering capabilities.

A smart toilet today is to microbiome diagnostics what a thermometer was to infectious disease in 1900: a useful instrument that detects something has changed, without being able to tell you what the something is. That is not a criticism. Thermometers save lives. But the reader who buys a smart toilet expecting a daily microbiome report will be disappointed.

Ingestible capsules: sampling from the inside

A more radical approach to real-time microbiome monitoring is to send the sensor to the microbiome, rather than waiting for the microbiome to come to the sensor. Ingestible capsule devices — pill-sized instruments that are swallowed, transit the gastrointestinal tract, and either collect samples or make measurements along the way — address a fundamental limitation of stool-based diagnostics.

That limitation is this: stool is a composite. By the time a bacterial community reaches the toilet, it has been mixed with contributions from every segment of the gut — stomach, duodenum, jejunum,

ileum, caecum, colon, rectum — and subjected to oxygen exposure that kills many of the strict anaerobes that dominate the living community. A stool sample tells you something about the gut microbiome, but it tells you most about the colonic microbiome and least about the small intestine, which is precisely where several important diseases — coeliac disease, small intestinal bacterial overgrowth, Crohn's ileitis — are located. If you want to know what is living in the jejunum, looking in the toilet is a poor strategy.

Ingestible capsules solve this by sampling directly. In 2025, a team at Tufts University published a small, pill-sized device with a 3D-printed soft elastic exterior whose sidewall inlets open in response to the pH change that occurs as the pill passes from the acidic stomach into the alkaline small intestine [REF:nema2025]. Elastic microvalves, using swellable polyacrylate beads, close the inlets once intestinal content has been collected, preserving the sample for analysis after the capsule is excreted. The device has completed preclinical characterisation, with human clinical trials planned.

A different engineering approach produced the CORAL capsule, which uses biomimetic structures — Triply Periodic Minimal Surfaces, inspired by marine coral geometry — to create a maze-like network of channels whose pore sizes are optimised to trap bacteria. The capsule captures a representative sample of the small-intestinal microbiome, which can then be analysed by standard sequencing after excretion.

Perhaps the most ambitious design combines biological sensors with electronics. A collaboration between MIT and Brigham and Women's Hospital produced an ingestible device containing engineered bacteria that function as living biosensors [REF:takahashi2024]. The bacteria are designed to detect thiosulfate — a marker of intestinal inflammation — and respond by producing bioluminescence. An onboard photodetector reads the light signal and transmits the data wirelessly to an external receiver. The device has been tested in pigs, where it successfully detected gastrointestinal inflammation in real time. A mask-sized version of a similar concept — the EBCare device, published in *Science* in 2025 — monitors exhaled breath condensate for biomarkers of respiratory inflammation and metabolic status using electrochemical biosensors and wireless readout [REF:ebcare2025].

And the technology is advancing rapidly. At ITF World 2025, the Dutch research centre imec presented an ingestible sensor prototype three times smaller than current capsule endoscopes — the first to measure redox balance in vivo, which provides a readout of microbial metabolic activity, alongside pH and temperature [REF:imec2025]. A 2026 review in *Advanced Robotics Research* described “multitasking smart intestinal capsule robots” capable of mucosal tissue and intestinal juice sampling, gas and pressure sensing, biomarker detection, and even targeted drug delivery — an in vivo diagnostic and therapeutic platform in a single swallowable pill.

These technologies are genuinely impressive. They are also genuinely early. All of them face engineering challenges that have not yet been solved: power (how do you run electronics and wireless transmission from a battery the size of a grain of rice?), biocompatibility (the capsule must not provoke an immune response or obstruct the intestine), data transmission (wireless signals from inside the abdomen are attenuated by tissue), and cost (a capsule that costs a thousand dollars to manufacture is not going to replace a ten-dollar stool test). The trajectory is clear — ingestible sensors will, eventually, allow clinicians to monitor the gut microbiome in vivo, in real time, and at specific locations. The timeline is not.

The continuous glucose monitor analogy

The most useful lens through which to view these emerging technologies is historical. Consider the trajectory of glucose monitoring for diabetes.

In the 1960s, blood glucose could only be measured in a hospital laboratory, from a venous blood sample, with a turnaround time of hours. In the 1980s, home glucose meters appeared — a finger-prick, a test strip, a result in minutes. In the 2000s, continuous glucose monitors (CGMs) arrived — a small sensor inserted under the skin, providing glucose readings every few minutes, transmitting data wirelessly to a phone. In the 2020s, CGMs are being worn by millions of non-diabetic people for general health tracking. The journey from laboratory to bedside to wristband took roughly fifty years.

Microbiome monitoring is, by this analogy, somewhere in the late 1960s. We can measure the microbiome in a laboratory, from a carefully collected and processed sample, with a turnaround time of days to weeks. The equivalent of the finger-prick glucose meter — a rapid, home-based microbiome test — does not yet exist, though some companies are working toward it. The equivalent of the CGM — continuous, real-time, in vivo monitoring — is the province of prototype capsules and academic publications. If the analogy holds, we might expect rapid point-of-care microbiome tests within ten to fifteen years and continuous monitoring within twenty to thirty. That is speculative, but it is not implausible.

30.4 AI and the Microbiome

The machine learning landscape

Everything in this chapter — every diagnostic classifier, every risk score, every biomarker panel — depends on machine learning. The microbiome produces data that is, by its nature, high-dimensional. A single stool sample may contain hundreds of bacterial species, thousands of genes, and tens of thousands of metabolic features. No human clinician can look at that data and intuit which features distinguish a cancer patient from a healthy control. Machine learning can.

The algorithms used for microbiome diagnostics are, for the most part, not exotic. Random forests — ensembles of decision trees that vote on the classification — are the workhorse of the field and have been since the early metagenomics era. Gradient-boosted methods (XGBoost, LightGBM) offer modest improvements. Neural networks have been applied but have not consistently outperformed simpler methods, partly because microbiome datasets are small by deep learning standards (hundreds to thousands of samples, where image recognition models are trained on millions).

The more interesting challenge is not which algorithm to use but how to make any algorithm work across populations. A 2025 study by Marcos-Zambrano and colleagues produced the most comprehensive assessment of this problem to date [REF:bestpractices2025]. They tested 156 combinations of preprocessing methods, batch-effect removal techniques, algorithms, and parameter settings across eighty-three case-control cohorts spanning twenty diseases, conducting nearly 13,000 internal and external validation experiments. From this massive benchmarking effort, they identified an optimal workflow: species-level taxonomic profiling, centred log-ratio transformation to handle the compositional nature of microbiome data, and random forest or gradient-boosted classification with careful cross-validation. Their recommended workflow demonstrated robust generalisability across diseases — but “robust” here means consistent, not perfect. The cross-cohort performance

drop that plagues every microbiome diagnostic study was reduced but not eliminated.

The generalisability problem deserves emphasis because it is the single biggest obstacle between the research literature and the clinic. Within a single study — where samples are collected, stored, and processed identically — microbiome classifiers often perform beautifully: AUC values of 0.85, 0.90, even 0.95 are common. But when the classifier trained on one study is tested on a different study — different hospital, different DNA extraction kit, different bioinformatics pipeline, different patient demographics — performance drops. The Parkinson's meta-analysis dropped from 0.72 to 0.61. CRC classifiers, among the most robust, still lose five to ten AUC points in cross-cohort validation. IBD classifiers show variable cross-ethnic performance.

The causes are well understood. Some are technical: different DNA extraction methods preferentially lyse different bacteria, changing the apparent composition of the sample. Different sequencing platforms have different error profiles. Different bioinformatics databases produce different taxonomic assignments. Some are biological: diet shapes the microbiome, and diet varies between populations. Genetics shape the microbiome, and genetic backgrounds vary between cohorts. The disease itself may have different microbial correlates in different populations — the same clinical diagnosis may arise through different ecological pathways.

None of these problems is unsolvable. Standardisation of collection, extraction, and sequencing protocols — the equivalent of the clinical chemistry standards that made blood tests reproducible across hospitals — would address the technical causes. Population-specific training sets, or the microbial-fingerprint normalisation approach developed by Ahn's group for melanoma, would address the biological causes. But standardisation requires consensus, investment, and time, and the field has not yet achieved it.

Foundation models: teaching AI to read microbial communities

A more ambitious approach to the generalisability problem comes from an unexpected direction: natural language processing.

In the past two years, researchers have begun to build what are called “foundation models” for the microbiome — large neural networks, modelled on the transformer architecture that powers ChatGPT and similar language models, that are pre-trained on vast amounts of unlabelled microbiome data before being fine-tuned for specific diagnostic tasks.

The intuition behind this approach is surprisingly natural. When a language model reads a sentence, it does not treat each word in isolation. It understands each word in the context of all the other words around it. The word “bank” means something different in “I sat on the river bank” and “I deposited money at the bank.” The model learns these contextual meanings not by being told explicitly but by reading billions of sentences and learning to predict which word comes next.

Microbial communities have the same property. *Akkermansia muciniphila* — a mucin-degrading bacterium we have met repeatedly in this book — means something different when it is surrounded by a diverse, fibre-fermenting community (where it is part of a healthy mucus-renewal cycle, and its gentle degradation of the mucus layer stimulates the goblet cells to produce more) and when it is the dominant organism in a depleted, post-antibiotic landscape (where it may be degrading the mucus layer faster than it can be replenished, contributing to barrier breakdown). A traditional machine learning model treats *Akkermansia* abundance as a single number. A foundation model can, in principle, represent *Akkermansia* differently depending on its community context — just as a language model represents “bank” differently depending on the surrounding sentence.

The largest microbiome foundation model published to date is the Microbial General Model (MGM), pre-trained on 263,302 microbiome samples using language modelling techniques [REF:mgm2024]. MGM treats each microbiome sample as a “sentence” in which the “words” are microbial taxa and their abundances. By learning to predict missing taxa and their abundances from the surrounding community, the model acquires contextualised representations that capture inter-species interactions and community-level patterns. When fine-tuned for specific diagnostic tasks, MGM showed significant improvements over traditional machine learning methods for microbial community classification.

A different approach, described in a 2023 *PLOS Computational Biology* paper, trained a deep language model on large-scale unlabelled microbiome data in a self-supervised fashion [REF:microbiome_lm2023]. The resulting model produced contextualised taxon representations in which a single microbial taxon is represented differently according to the specific microbial environment in which it appears — exactly the kind of context-dependent interpretation that traditional methods lack.

The promise of these models is that they may solve the generalisability problem by learning to distinguish between “microbiome features that are universally diagnostic” and “microbiome features that are population-specific background variation.” If MGM has seen 263,000 samples from dozens of populations, it has, in principle, learned what a “normal” microbiome looks like in each of those populations. When it encounters a new sample from a new population, it can normalise for the population-specific background and focus on the disease-specific signal. This is the same logic as the microbial-fingerprint approach, but implemented at a deeper level.

The honest assessment: foundation models for the microbiome are in their infancy. The 263,000 samples that trained MGM are a large number by microbiome standards but minuscule compared to the datasets that power modern language models. Performance improvements over traditional methods are real but modest. The interpretability problem — understanding why a model classifies a sample as “disease” — is at least as severe in microbiome AI as in other domains. Clinical regulators are unlikely to approve a diagnostic test that provides accurate predictions but cannot explain them. The path from impressive benchmark performance to an approved clinical diagnostic is long, and most AI-based diagnostic tools in other fields — radiology, pathology, dermatology — have taken five to ten years to traverse it. But the trajectory is unmistakable, and the reader should expect to hear much more about microbiome foundation models in the decade to come.

Personalised nutrition: where the algorithm meets the consumer

The diagnostic platform that has, so far, reached the most consumers is not a disease test at all. It is the personalised nutrition algorithm — a system that uses microbiome data, alongside other inputs, to predict how an individual will respond to specific foods and to generate personalised dietary recommendations.

The scientific foundation was laid in 2015 by Eran Segal and Eran Elinav at the Weizmann Institute of Science [REF:zeevi2015]. Their study — which we touched on in Chapter 27 — continuously monitored blood glucose levels in 800 people over a week, measured responses to 46,898 meals, and discovered that individuals varied enormously in their glycaemic response to identical foods. One person’s blood sugar spiked after eating a banana but not after a cookie; for another person, the pattern was reversed. A machine learning algorithm integrating blood parameters, dietary habits, body measurements, physical activity, and — importantly — gut microbiome composition could predict these personalised responses with remarkable accuracy. When the algorithm was used to

design “good” and “bad” diets for individual participants, the algorithmically designed diets produced significantly lower blood sugar spikes and, intriguingly, consistent shifts in the participants’ gut microbiome compositions.

This study has since spawned two major commercial platforms. DayTwo, which licensed the Weizmann algorithm, provides microbiome-based dietary recommendations for glycaemic control and has been used by approximately 75,000 people, in partnership with the Mayo Clinic and Johnson & Johnson. ZOE, a UK-based company, offers a broader personalised nutrition platform that combines continuous glucose monitoring, blood lipid responses, and gut microbiome profiling. ZOE’s METHOD trial — a randomised controlled trial published in 2024 — showed that personalised dietary advice based on their algorithm produced significant reductions in triglycerides, body weight (minus 2.5 kilograms), and waist circumference (minus 2.4 centimetres) compared with standard dietary guidance, along with increased gut microbiome diversity.

These platforms occupy a curious regulatory space. They are not marketed as medical devices — which would require FDA approval — but they use diagnostic data (microbiome composition, continuous glucose monitoring) to make health recommendations. The consumer receives what looks and feels like a medical test result — “your gut microbiome diversity is in the 30th percentile” — without the regulatory oversight that a medical test would carry. This is the territory that the Lancet consensus statement is addressing, and it connects directly to the commercial landscape we turn to next.

30.5 The Direct-to-Consumer Problem

The commercial microbiome testing industry is booming. As of mid-2025, at least thirty companies sell direct-to-consumer microbiome testing kits in the developed world. The global market is valued at approximately \$1.2 billion and is projected to reach \$3.5 billion by 2033. The consumer buys a kit, collects a stool sample at home, mails it to the company, and receives — within a few weeks — a report that may include a list of the bacteria detected in their gut, a “diversity score,” diet and supplement recommendations, and sometimes a disease risk assessment.

The problem is that these reports, in most cases, are not supported by the kind of evidence that this chapter has described. The studies we have reviewed — the Wirbel and Thomas meta-analyses for CRC, the Zheng study for IBD, the Menozzi study for Parkinson’s — were conducted with rigorous protocols, multi-centre validation, and statistical methods designed to detect and correct for the many biases that microbiome data is prone to. The typical DTC microbiome test shares almost none of these safeguards.

Several specific problems recur. First, reproducibility: the same stool sample, split and sent to two different DTC providers, may produce contradictory results — different species lists, different diversity scores, different dietary recommendations. This is because the companies use different DNA extraction methods, different sequencing platforms, different bioinformatics databases, and different proprietary algorithms. There is no standardisation. Second, clinical validation: most DTC companies have not published peer-reviewed studies demonstrating that their specific test, on their specific platform, predicts any clinically meaningful outcome. They reference the broader microbiome literature — the same literature we have reviewed in this chapter — but the leap from “researchers at EMBL found that *Fusobacterium nucleatum* is enriched in CRC” to “our test can tell you whether you should eat more fibre” is vast and unsupported. Third, interpretation: the reports

are often presented in a format that implies clinical authority — colour-coded risk scores, comparisons to “healthy” reference ranges — but there is no consensus on what a “healthy” microbiome looks like (a point we made in Chapter 1), and the variability of the microbiome from day to day and meal to meal means that a single snapshot may not be representative.

The most dramatic failure in the DTC space was uBiome, once a market leader. In 2019, the FBI raided the company’s offices. Federal prosecutors alleged that uBiome had engaged in systematic insurance fraud: pressuring consumers through repetitive emails to submit multiple stool samples for medically unnecessary testing, then billing private health insurers and Medicare two or three times for the same set of tests. The company collapsed. Its founders faced criminal charges. The episode was a stark illustration of what happens when commercial incentives run far ahead of scientific evidence.

A 2024 editorial in *Science* by a group of leading microbiome researchers argued that the DTC testing industry needs more regulation [REF:dtc_science2024]. A companion editorial in the *Lancet Gastroenterology & Hepatology* made the same case [REF:dtc_lancet2024]. And a 2025 analysis in the *Journal of Law and the Biosciences* examined whether the current regulatory framework is sufficient to protect consumers from medical, economic, and dignitary harms — and concluded that it is not [REF:dtc_jlb2025].

The regulatory landscape is evolving, but not in a reassuring direction. In September 2025, the FDA rescinded a 2024 rule that would have classified laboratory-developed tests (LDTs) — the category into which most DTC microbiome tests fall — as medical devices subject to FDA oversight. This decision leaves DTC microbiome tests in a regulatory grey zone: they are neither approved nor prohibited, and the consumer has no federal assurance that the test they are buying produces meaningful results. The Lancet consensus statement — while careful not to condemn all microbiome testing — provides a set of minimum requirements that no current DTC product fully meets.

This is not to say that all commercial microbiome testing is useless. ZOE publishes peer-reviewed RCTs. DayTwo is built on a well-validated algorithm licensed from a world-class research institution. Some companies, including those we have not named, are investing in clinical validation studies. But the industry as a whole is ahead of its evidence base, and the consumer who buys a microbiome test expecting the kind of diagnostic precision described in the first half of this chapter is likely to be disappointed.

The reader who finishes this section should be equipped with a simple checklist for evaluating any microbiome test, commercial or clinical:

Has the test been validated in multiple independent cohorts? If not, the results may not be reproducible in your population. Has the test been shown to predict a clinically meaningful outcome (disease, treatment response, health improvement), or does it merely report microbiome composition? If the latter, the information may be interesting but is not diagnostic. Is the company publishing its methods and validation data in peer-reviewed journals, or relying on proprietary claims? Transparency is not a guarantee of quality, but its absence is a warning sign. Does the test’s performance exceed what could be achieved with simpler, cheaper methods — a dietary questionnaire, a standard blood test, a clinical history? If not, the microbiome test may be adding complexity without adding value.

30.6 What Is Ready, What Is Not, and What Is Coming

This chapter has covered a great deal of ground, from stool-based cancer screening to ingestible sensor pills, from machine learning classifiers to the wreckage of retracted papers. A calibrated summary is in order.

What is close to clinical use. Stool-based microbiome diagnostics for colorectal cancer and inflammatory bowel disease have reached the stage of large, multi-cohort validation with cross-ethnic replication. The IBD diagnostic developed by Zheng and colleagues has been translated into a PCR-based format that any diagnostic laboratory could run. These tests are not yet approved, but the evidence supporting them is comparable in quality to the evidence that supported other diagnostic innovations — such as the faecal immunochemical test itself — at a similar stage in their development. If prospective clinical utility trials confirm that microbiome-based screening improves patient outcomes, regulatory approval could follow within five to ten years.

What is five to ten years away. Companion diagnostics for immunotherapy response are promising but require independent, multi-centre, prospective validation before they can be recommended. The pharmacomicrobiomics vision — pre-treatment microbiome profiling to predict drug response across a range of medications — is scientifically grounded but has produced no clinically validated test outside immunotherapy. Personalised nutrition platforms based on microbiome data have reached the consumer market ahead of regulatory oversight, and their clinical utility — while supported by some RCT evidence — remains debated.

What is ten to twenty years away. Real-time microbiome monitoring via ingestible capsules, smart toilets with actual microbial profiling capability, and continuous metabolite sensing via wearable devices are in the engineering prototype stage. The trajectory mirrors that of continuous glucose monitoring, but the technical challenges are greater: the microbiome is orders of magnitude more complex than a single blood glucose measurement, and miniaturising DNA sequencing or metabolite detection to the scale of an ingestible pill remains formidable. Foundation models for microbiome data — the AI equivalent of learning to read microbial communities as a language — are in their earliest stages but offer a plausible path to solving the generalisability problem that limits all current diagnostics.

What the reader should take away. The microbiome is not yet a clinical diagnostic tool, but it is no longer merely a research curiosity. The evidence supporting microbiome-based diagnostics in several specific applications — CRC, IBD, liver disease — is now stronger than the evidence that supported many now-standard clinical tests at the same stage of their development. The barriers that remain are not scientific in the sense that we lack knowledge; they are translational in the sense that we lack the standardised protocols, prospective trials, and regulatory frameworks needed to move from research to practice. Those barriers will be overcome. The question is not whether the microbiome will enter clinical diagnostics, but when — and, critically, whether it will enter through the front door of rigorous validation or through the back door of unregulated commercial products.

The next chapter turns to the ethical and social questions that this diagnostic promise raises. If a stool test can tell you that your microbiome resembles an early-Parkinson's microbiome, should you be told? If a microbiome report influences your insurance premiums, who should have access to the data? If the benefits of microbiome medicine are available only to those who can afford a \$200 test and a personalised nutrition subscription, does the microbiome revolution widen or narrow health disparities? These are not future questions. They are present ones, and they deserve the same rigorous treatment that the science itself demands.

Key References

- [REF:icc2024] International consensus statement on microbiome testing in clinical practice. *Lancet Gastroenterology & Hepatology*, 2025;10(2):154–167. DOI: 10.1016/S2468-1253(24)00311-X
- [REF:zeller2014] Zeller, G., Tap, J., Voigt, A.Y., Sunagawa, S., Kultima, J.R., Costea, P.I., et al. (2014). Potential of fecal microbiota for early-stage detection of colorectal cancer. *Molecular Systems Biology*, 10(11), 766. DOI: 10.15252/msb.20145645
- [REF:wirbel2019] Wirbel, J., Pyl, P.T., Kartal, E., Zych, K., Kashani, A., Milanese, A., et al. (2019). Meta-analysis of fecal metagenomes reveals global microbial signatures that are specific for colorectal cancer. *Nature Medicine*, 25(4), 679–689. DOI: 10.1038/s41591-019-0406-6
- [REF:thomas2019] Thomas, A.M., Manghi, P., Asnicar, F., Pasolli, E., Armanini, F., Zolfo, M., et al. (2019). Metagenomic analysis of colorectal cancer datasets identifies cross-cohort microbial diagnostic signatures and a link with choline degradation. *Nature Medicine*, 25(4), 667–678. DOI: 10.1038/s41591-019-0405-7
- [REF:pooled2025] Pooled analysis of 3,741 stool metagenomes from 18 cohorts for cross-stage and strain-level reproducible microbial biomarkers of colorectal cancer. *Nature Medicine*, 2025. DOI: 10.1038/s41591-025-03693-9
- [REF:trickovic2025] Tričković, M., Kieser, S., Zdobnov, E.M. & Trajkovski, M. (2025). Subspecies-level resolution of the human gut microbiome for colorectal cancer detection. *Cell Host & Microbe*, 33(8). DOI: 10.1016/j.chom.2025.07.015
- [REF:zheng2024] Zheng, J., Sun, Q., Zhang, M., et al. (2024). Noninvasive, microbiome-based diagnosis of inflammatory bowel disease. *Nature Medicine*, 30(12), 3555–3567. DOI: 10.1038/s41591-024-03280-4
- [REF:caussy2019] Caussy, C., Tripathi, A., Humphrey, G., et al. (2019). A gut microbiome signature for cirrhosis due to nonalcoholic fatty liver disease. *Nature Communications*, 10(1), 1406. DOI: 10.1038/s41467-019-09455-9
- [REF:iebbba2020] Iebba, V., et al. (2020). A universal gut microbiome-derived signature predicts cirrhosis. *Cell Metabolism*, 32(5), 878–888. DOI: 10.1016/j.cmet.2020.06.005
- [REF:menozzi2026] Menozzi, E., Schapira, A.H.V., et al. (2026). Microbiome signature of Parkinson's disease in healthy and genetically at-risk individuals. *Nature Medicine*. DOI: 10.1038/s41591-026-04318-5
- [REF:romano2025] Romano, S., et al. (2025). Machine learning-based meta-analysis reveals gut microbiome alterations associated with Parkinson's disease. *Nature Communications*. DOI: 10.1038/s41467-025-56829-3
- [REF:oralage2026] Oral microbiome signatures predict biological age and host health. *Nature Communications*, 2026. DOI: 10.1038/s41467-026-72096-2
- [REF:oralpc2025] Oral microbiome as a diagnostic biomarker for pancreatic cancer: a systematic review and meta-analysis of diagnostic accuracy. *Acta Odontologica Scandinavica*,

2025. DOI: 10.1080/20002297.2025.2571188

- [REF:poore2020_RETRACTED] Poore, G.D., Kopylova, E., Zhu, Q., et al. (2020). Microbiome analyses of blood and tissues suggest cancer diagnostic approach. *Nature*, 579(7800), 567–574. DOI: 10.1038/s41586-020-2095-1 **(RETRACTED June 2024)**
- [REF:gihaw2023] Gihawi, A., Ge, Y., Lu, J., et al. (2023). Major data analysis errors invalidate cancer microbiome findings. *Nature*. DOI: 10.1038/s41586-023-06593-1
- [REF:liang2024] Liang, W., et al. (2024). Circulating microbiome DNA as biomarkers for early diagnosis and recurrence of lung cancer. *Cell Reports Medicine*, 5(4), 101480. DOI: 10.1016/S2666-3791(24)00167-8
- [REF:biomeone2023] Assessing the Performance of a Novel Stool-Based Microbiome Test That Predicts Response to First Line Immune Checkpoint Inhibitors in Multiple Cancer Types. *Diagnostics*, 2023;13(14):2329. DOI: 10.3390/diagnostics13142329
- [REF:blautia2025] Non-intestinal microbial signatures in stool as predictors of cancer immunotherapy outcome. *bioRxiv*, 2025. DOI: 10.1101/2025.05.07.652660
- [REF:usyk2026] Usyk, M., Li, H., Osman, I., Weber, J.S., Knight, R., González, A., Hayes, R.B. & Ahn, J. (2026). Gut microbiome is associated with recurrence-free survival in patients with resected high-risk melanoma receiving adjuvant immune checkpoint blockade. *Cell*. DOI: 10.1016/j.cell.2026.03.041
- [REF:bullman2024] Bullman, S., et al. (2024). A distinct *Fusobacterium nucleatum* clade dominates the colorectal cancer niche. *Nature*, 628, 424–432. DOI: 10.1038/s41586-024-07182-w
- [REF:park2020] Park, S.-M., Won, D.D., Lee, B.J., et al. (2020). A mountable toilet system for personalized health monitoring via the analysis of excreta. *Nature Biomedical Engineering*, 4, 624–635. DOI: 10.1038/s41551-020-0534-9
- [REF:nema2025] Nema, V., et al. (2025). A small pill-like ingestible microdevice for site-specific microbiome sampling in the upper GI tract. *Small*. DOI: 10.1002/smll.202510289
- [REF:takahashi2024] MIT/BWH collaboration (2024). Ingestible device combining engineered bacteria with electronics for real-time gut inflammation monitoring. *[Published in Device/Nature Biomedical Engineering]*
- [REF:imec2025] imec/OnePlanet Research Center (2025). Ingestible sensor prototype for redox, pH, and temperature monitoring. Presented at ITF World 2025.
- [REF:ebcare2025] EBCare: A smart mask for exhaled breath condensate harvesting and analysis. *Science*, 2025. DOI: 10.1126/science.adn6471
- [REF:bestpractices2025] Best practices for developing microbiome-based disease diagnostic classifiers through machine learning. *Gut Microbes*, 2025;17(1):2489074. DOI: 10.1080/19490976.2025.2489074
- [REF:mgm2024] Microbial General Model (MGM): a large-scale pretrained foundation model for microbiome analyses. *bioRxiv*, 2024. DOI: 10.1101/2024.12.30.630825
- [REF:microbiome_lm2023] Learning a deep language model for microbiomes: the power of large-scale unlabeled microbiome data. *PLOS Computational Biology*, 2023;19(6):e1011353.

DOI: 10.1371/journal.pcbi.1011353

- [REF:zeevi2015] Zeevi, D., Korem, T., Zmora, N., et al. (2015). Personalized nutrition by prediction of glycemic responses. *Cell*, 163(5), 1079–1094. DOI: 10.1016/j.cell.2015.11.001
- [REF:dtc_science2024] The DTC microbiome testing industry needs more regulation. *Science*, 2024. DOI: 10.1126/science.adk4271
- [REF:dtc_lancet2024] Direct-to-consumer microbiome testing needs regulation. *Lancet Gastroenterology & Hepatology*, 2024;9(8):695. DOI: 10.1016/S2468-1253(24)00163-8
- [REF:dtc_jlb2025] Is the current regulatory framework for direct-to-consumer microbiome-based tests sufficient to protect consumers from medical, economic, and dignitary harms? *Journal of Law and the Biosciences*, 2025;12(2):lsafo24. DOI: 10.1093/jlb/lsafo24





The Inhabited Body — Bibliography by Chapter

Curated from Paperpile library (185 references) April 2026

References sorted by primary relevance to each chapter. Many references appear in multiple chapters; listed here under their *primary* home. Diabetes-focused references are concentrated in Chapters 22, 26, and 27 where they are most relevant, rather than scattered throughout.

Ch. 1 — You Are Never Alone

- [REF:sender2016] Sender R, Fuchs S & Milo R (2016). Revised estimates of the number of human and bacteria cells in the body. *PLOS Biology*, 14(8), e1002533. doi:10.1371/journal.pbio.1002533
- [REF:luckey1972] Luckey TD (1972). Introduction to intestinal microecology. *Am. J. Clin. Nutr.*, 25(12), 1292–1294.
- [REF:berg1996] Berg RD (1996). The indigenous gastrointestinal microflora. *Trends Microbiol.*, 4(11), 430–435. doi:10.1016/0966-842X(96)10057-3
- [REF:bianconi2013] Bianconi E et al. (2013). An estimation of the number of cells in the human body. *Ann. Hum. Biol.*, 40(6), 463–471. doi:10.3109/03014460.2013.807878
- [REF:berg2020] Berg G et al. (2020). Microbiome definition re-visited: old concepts and new challenges. *Microbiome*, 8, 103. doi:10.1186/s40168-020-00875-0
- [REF:gilbert2018] Gilbert JA et al. (2018). Current understanding of the human microbiome. *Nat. Med.*, 24(4), 392–400. doi:10.1038/nm.4517
- [REF:qin2010] Qin J et al. (2010). A human gut microbial gene catalogue established by metagenomic sequencing. *Nature*, 464(7285), 59–65. doi:10.1038/nature08821
- [REF:tierney2019] Tierney BT et al. (2019). The landscape of genetic content in the gut and oral human microbiome. *Cell Host Microbe*, 26(2), 283–295. doi:10.1016/j.chom.2019.07.008
- [REF:almeida2021] Almeida A et al. (2021). A unified catalog of 204,938 reference genomes from the human gut microbiome. *Nat. Biotechnol.*, 39, 105–114. doi:10.1038/s41587-020-0603-3

Ch. 2 — Built Together: How Microbes Shaped Human Evolution

- Mao J et al. (2021). Human endogenous retroviruses in development and disease. *Comput. Struct. Biotechnol. J.* doi:10.1016/j.csbj.2021.10.037
- Mortelmans K et al. (2016). The role of human endogenous retroviruses in brain development and function. *APMIS*. doi:10.1111/apm.12495
- Mietzsch M & Agbandje-McKenna M (2017). The Good That Viruses Do. *Annu Rev Virol.* doi:10.1146/annurev-vi-04-071217-100011
- Roossinck MJ (2015). Move over, bacteria! Viruses make their mark as mutualistic microbial symbionts. *J. Virol.* doi:10.1128/JVI.02974-14
- American Society for Microbiology (2015). Viruses: You’ve heard the bad; here’s the good. *Science Daily*.

Ch. 3 — Your First Inheritance: Colonisation Across a Lifetime

- Argaw-Denboba A et al. (2024). Paternal microbiome perturbations impact offspring fitness. *Nature*. doi:10.1038/s41586-024-07336-w
- Veerus L et al. (2024). Dad's gut microbes matter for pregnancy health and baby's growth. *Nature*. doi:10.1038/d41586-024-01191-5
- Way SS (2026). Some of your cells are not genetically yours — what can they tell us about life and death? *Nature*. doi:10.1038/d41586-025-04102-4
- Soderborg TK et al. (2016). Microbial transmission from mothers with obesity or diabetes to infants. *Diabetologia*. doi:10.1007/s00125-016-3880-0
- Beghetti I et al. (2019). Human Milk's Hidden Gift: Implications of the Milk Microbiome for Preterm Infants' Health. *Nutrients*. doi:10.3390/nu11122944
- Hinde K & Lewis ZT (2015). Mother's littlest helpers. *Science*. doi:10.1126/science.aac7436
- LeMay-Nedjelski L et al. (2021). Maternal Diet and Infant Feeding Practices Are Associated with Variation in the Human Milk Microbiota. *J. Nutr.* doi:10.1093/jn/nxaa248
- Beghetti I et al. (2022). Donor Human Milk Microbiota: Comparison with Preterm HM and the Effect of Pasteurization. *Nutrients*. doi:10.3390/nu14122483
- Collado MC & Segata N (2020). Initial exploration of in utero microbial colonization. *Nat. Med.* doi:10.1038/s41591-020-0836-1
- Chavarro JE et al. (2020). Association of Birth by Cesarean Delivery With Obesity and T2D Among Adult Women. *JAMA Netw Open*. doi:10.1001/jamanetworkopen.2020.2605
- Shen X et al. (2024). Nonlinear dynamics of multi-omics profiles during human aging. *Nat. Aging*. doi:10.1038/s43587-024-00692-2
- Zhuang C et al. (2021). Association between antibiotics during pregnancy and obesity in 5-year-old children. *Transl Pediatr*. doi:10.21037/tp-21-220

Ch. 4 — Tools of Discovery

- Jones MB et al. (2015). Library preparation methodology can influence genomic and functional predictions in human microbiome research. *PNAS*. doi:10.1073/pnas.1519288112
- Wang J & Jia H (2016). Metagenome-wide association studies: fine-mining the microbiome. *Nat. Rev. Microbiol.* doi:10.1038/nrmicro.2016.83
- Kishikawa T et al. (2022). OMARU: a robust pipeline for metagenome-wide association study. *NAR Genom Bioinform.* doi:10.1093/nargab/lqac019
- D'Andreano S et al. (2021). Rapid real-time identification of fungi with long amplicon nanopore sequencing. *Biol Methods Protoc.* doi:10.1093/biomethods/bpaa026
- Mitsuhashi S et al. (2017). Portable system for rapid bacterial composition analysis using nanopore sequencer. *Sci. Rep.* doi:10.1038/s41598-017-05772-5
- Taxt AM et al. (2020). Rapid identification of pathogens, ARGs and plasmids in blood cultures by nanopore sequencing. *Sci. Rep.* doi:10.1038/s41598-020-64616-x

Ch. 5 — The Gut

- [REF:tropini2017] Tropini C, Earle KA, Huang KC & Sonnenburg JL (2017). The Gut Microbiome: Connecting Spatial Organization to Function. *Cell Host Microbe*, 21(4), 433–442. doi:10.1016/j.chom.2017.03.010
- [REF:johansson2008] Johansson MEV, Phillipson M, Petersson J, Velcich A, Holm L & Hansson GC (2008). The inner of the two Muc2 mucin-dependent mucus layers in colon is devoid of bacteria. *PNAS*, 105(39), 15064–15069. doi:10.1073/pnas.0803124105

- [REF:desai2016] Desai MS et al. (2016). A Dietary Fiber-Deprived Gut Microbiota Degrades the Colonic Mucus Barrier and Enhances Pathogen Susceptibility. *Cell*, 167(5), 1339–1353. doi:10.1016/j.cell.2016.10.043
- [REF:birchenough2016] Birchenough GMH, Nyström EEL, Johansson MEV & Hansson GC (2016). A sentinel goblet cell guards the colonic crypt by triggering Nlrp6-dependent Muc2 secretion. *Science*, 352(6293), 1535–1542. doi:10.1126/science.aaf7419
- [REF:hansson2020] Hansson GC (2020). Mucins and the Microbiome. *Annu. Rev. Biochem.*, 89, 769–793. doi:10.1146/annurev-biochem-011520-105053
- [REF:cummings1987] Cummings JH, Pomare EW, Branch WJ, Naylor CP & Macfarlane GT (1987). Short chain fatty acids in human large intestine, portal, hepatic and venous blood. *Gut*, 28(10), 1221–1227. doi:10.1136/gut.28.10.1221
- [REF:koh2016] Koh A, De Vadder F, Kovatcheva-Datchary P & Bäckhed F (2016). From Dietary Fiber to Host Physiology: Short-Chain Fatty Acids as Key Bacterial Metabolites. *Cell*, 165(6), 1332–1345. doi:10.1016/j.cell.2016.05.041
- [REF:morrison2016] Morrison DJ & Preston T (2016). Formation of short chain fatty acids by the gut microbiota and their impact on human metabolism. *Gut Microbes*, 7(3), 189–200. doi:10.1080/19490976.2015.1134082
- [REF:furness2012] Furness JB (2012). The enteric nervous system and neurogastroenterology. *Nat. Rev. Gastroenterol. Hepatol.*, 9(5), 286–294. doi:10.1038/nrgastro.2012.32
- [REF:yano2015] Yano JM et al. (2015). Indigenous bacteria from the gut microbiota regulate host serotonin biosynthesis. *Cell*, 161(2), 264–276. doi:10.1016/j.cell.2015.02.047
- [REF:arumugam2011] Arumugam M et al. (2011). Enterotypes of the human gut microbiome. *Nature*, 473(7346), 174–180. doi:10.1038/nature09944
- [REF:knights2014] Knights D et al. (2014). Rethinking “enterotypes.” *Cell Host Microbe*, 16(4), 433–437. doi:10.1016/j.chom.2014.09.013
- [REF:costea2018] Costea PI et al. (2018). Enterotypes in the landscape of gut microbial community composition. *Nat. Microbiol.*, 3(1), 8–16. doi:10.1038/s41564-017-0072-8
- [REF:wu2011] Wu GD et al. (2011). Linking long-term dietary patterns with gut microbial enterotypes. *Science*, 334(6052), 105–108. doi:10.1126/science.1208344
- [REF:rinninella2019] Rinninella E et al. (2019). What is the Healthy Gut Microbiota Composition? A Changing Ecosystem across Age, Environment, Diet, and Diseases. *Microorganisms*, 7(1), 14. doi:10.3390/microorganisms7010014
- [REF:bohnhoff1954] Bohnhoff M, Drake BL & Miller CP (1954). Effect of streptomycin on susceptibility of intestinal tract to experimental Salmonella infection. *Proc. Soc. Exp. Biol. Med.*, 86(1), 132–137.
- [REF:vanderwaaij1971] van der Waaij D, Berghuis-de Vries JM & Lekkerkerk-van der Wees JEC (1971). Colonization resistance of the digestive tract in conventional and antibiotic-treated mice. *J. Hygiene*, 69(3), 405–411. doi:10.1017/S0022172400021653
- [REF:buffie2013] Buffie CG & Pamer EG (2013). Microbiota-mediated colonization resistance against intestinal pathogens. *Nat. Rev. Immunol.*, 13(11), 790–801. doi:10.1038/nri3535
- [REF:ducarmon2019] Ducarmon QR et al. (2019). Gut Microbiota and Colonization Resistance against Bacterial Enteric Infection. *Microbiol. Mol. Biol. Rev.*, 83(3), e00007-19. doi:10.1128/MMBR.00007-19
- Mallott EK et al. (2024). The primate gut microbiota contributes to interspecific differences in host metabolism. *Microb. Genom.* doi:10.1099/mgen.0.001322
- Corbin KD et al. (2023). Host-diet-gut microbiome interactions influence human energy balance. *Nat. Commun.* doi:10.1038/s41467-023-38778-x
- Collins J et al. (2018). Dietary trehalose enhances virulence of epidemic *Clostridium difficile*.

Nature. doi:10.1038/nature25178

- Wikberg EC et al. (2020). Interactions between social groups of colobus monkeys explain similarities in their gut microbiomes. *Anim. Behav.* doi:10.1016/j.anbehav.2020.02.011
- Odenwald MA & Turner JR (2017). The intestinal epithelial barrier: a therapeutic target? *Nat. Rev. Gastroenterol. Hepatol.* doi:10.1038/nrgastro.2016.169
- Mu Q et al. (2017). Leaky Gut As a Danger Signal for Autoimmune Diseases. *Front. Immunol.* doi:10.3389/fimmu.2017.00598

Ch. 6 — The Skin

- [REF:grice2009] Grice EA, Kong HH, Conlan S, Deming CB, Davis J, Young AC, ... & Segre JA (2009). Topographical and temporal diversity of the human skin microbiome. *Science*, 324(5931), 1190–1192. doi:10.1126/science.1171700
- [REF:gricesegre2011] Grice EA & Segre JA (2011). The skin microbiome. *Nat. Rev. Microbiol.*, 9(4), 244–253. doi:10.1038/nrmicro2537
- [REF:byrd2018] Byrd AL, Belkaid Y & Segre JA (2018). The human skin microbiome. *Nat. Rev. Microbiol.*, 16(3), 143–155. doi:10.1038/nrmicro.2017.157
- [REF:findley2013] Findley K, Oh J, Yang J, Conlan S, Deming C, Meyer JA, ... & Segre JA (2013). Topographic diversity of fungal and bacterial communities in human skin. *Nature*, 498(7454), 367–370. doi:10.1038/nature12171
- [REF:kong2012] Kong HH, Oh J, Deming C, Conlan S, Grice EA, Beatson MA, ... & Segre JA (2012). Temporal shifts in the skin microbiome associated with disease flares and treatment in children with atopic dermatitis. *Genome Res.*, 22(5), 850–859. doi:10.1101/gr.131029.111
- [REF:strachan1989] Strachan DP (1989). Hay fever, hygiene, and household size. *BMJ*, 299(6710), 1259–1260. doi:10.1136/bmj.299.6710.1259
- [REF:rook2014] Rook GAW, Raison CL & Lowry CA (2014). Microbiota, immunoregulatory old friends and psychiatric disorders. *Adv. Exp. Med. Biol.*, 817, 319–356. doi:10.1007/978-1-4939-0897-4_15
- [REF:pfefferle2021] Pfefferle PI, Keber CU, Cohen RM & Garn H (2021). The Hygiene Hypothesis — Learning From but Not Living in the Past. *Front. Immunol.*, 12, 635935. doi:10.3389/fimmu.2021.635935
- [REF:nakatsuji2017] Nakatsuji T, Chen TH, Narala S, Chun KA, Two AM, Yun T, ... & Gallo RL (2017). Antimicrobials from human skin commensal bacteria protect against *Staphylococcus aureus* and are deficient in atopic dermatitis. *Sci. Transl. Med.*, 9(378), eaah4680. doi:10.1126/scitranslmed.aah4680

Ch. 7 — The Respiratory Tract

- [REF:hilty2010] Hilty M, Burke C, Pedro H, Cardenas P, Bush A, Bossley C, ... & Cookson WOC (2010). Disordered microbial communities in asthmatic airways. *PLoS One*, 5(1), e8578. doi:10.1371/journal.pone.0008578
- [REF:dickson2015a] Dickson RP, Erb-Downward JR & Huffnagle GB (2015). Homeostasis and its disruption in the lung microbiome. *Am. J. Physiol. Lung Cell. Mol. Physiol.*, 309(10), L1047–L1055. doi:10.1152/ajplung.00279.2015
- [REF:dickson2016] Dickson RP, Erb-Downward JR, Martinez FJ & Huffnagle GB (2016). The microbiome and the respiratory tract. *Annu. Rev. Physiol.*, 78, 481–504. doi:10.1146/annurev-physiol-021115-105238
- [REF:dickson2015b] Dickson RP (2016). The microbiome and critical illness. *Lancet Respir.*

Med., 4(1), 59–72. doi:10.1016/S2213-2600(15)00427-0

- [REF:segal2013] Segal LN, Alekseyenko AV, Clemente JC, Kulkarni R, Wu B, Gao Z, ... & Weiden MD (2013). Enrichment of lung microbiome with supraglottic taxa is associated with increased pulmonary inflammation. *Microbiome*, 1(1), 19. doi:10.1186/2049-2618-1-19
- [REF:segal2016] Segal LN, Clemente JC, Tsay JCJ, Koralov SB, Keller BC, Wu BG, ... & Weiden MD (2016). Enrichment of the lung microbiome with oral taxa is associated with lung inflammation of a Th17 phenotype. *Nat. Microbiol.*, 1, 16031. doi:10.1038/nmicrobiol.2016.31
- [REF:cleary2017] Cleary DW & Clarke SC (2017). The nasopharyngeal microbiome. *Emerg. Top. Life Sci.*, 1(4), 297–312. doi:10.1042/ETLS20170041
- [REF:desteenhuijsen2016] de Steenhuijsen Pitsers WAA, Heinonen S, Hasrat R, Bunsow E, Smith B, Suarez-Arrabal MC, ... & Mejias A (2016). Nasopharyngeal microbiota, host transcriptome, and disease severity in children with respiratory syncytial virus infection. *Am. J. Respir. Crit. Care Med.*, 194(9), 1104–1115. doi:10.1164/rccm.201602-0220OC
- [REF:mammen2016] Mammen MJ & Sethi S (2016). COPD and the microbiome. *Respirology*, 21(4), 590–599. doi:10.1111/resp.12732
- [REF:han2012] Han MK, Huang YJ, Lipuma JJ, Boushey HA, Boucher RC, Cookson WO, ... & Martinez FJ (2012). Significance of the microbiome in obstructive lung disease. *Thorax*, 67(5), 456–463. doi:10.1136/thoraxjnl-2011-201183
- [REF:chung2017] Chung KF (2017). Airway microbial dysbiosis in asthmatic patients: A target for prevention and treatment? *J. Allergy Clin. Immunol.*, 139(4), 1071–1081. doi:10.1016/j.jaci.2017.02.004
- [REF:chung2017ats] Chung KF (2017). Potential role of the lung microbiome in shaping asthma phenotypes. *Ann. Am. Thorac. Soc.*, 14(Suppl 5), S326–S331. doi:10.1513/AnnalsATS.201702-138AW
- [REF:dicicco2018] Di Cicco M, Pistello M, Jacinto T, Ragazzo V, Piras M, Freer G, Pifferi M & Peroni D (2018). Does lung microbiome play a causal or casual role in asthma? *Pediatr. Pulmonol.*, 53(10), 1340–1345. doi:10.1002/ppul.24086
- [REF:frayman2017] Frayman KB, Armstrong DS, Grimwood K & Ranganathan SC (2017). The airway microbiota in early cystic fibrosis lung disease. *Pediatr. Pulmonol.*, 52(11), 1384–1404. doi:10.1002/ppul.23782
- [REF:cuthbertson2020] Cuthbertson L, Felton I, James P, Cox MJ, Bilton D, Schelenz S, ... & Moffatt MF (2020). The fungal airway microbiome in cystic fibrosis and non-cystic fibrosis bronchiectasis. *J. Cyst. Fibros.*, 20(2), 295–302. doi:10.1016/j.jcf.2020.05.013
- [REF:jenks2021] Jenks JD, Nam HH & Hoenigl M (2021). Invasive aspergillosis in critically ill patients: Review of definitions and diagnostic approaches. *Mycoses*, 64(9), 1002–1014. doi:10.1111/myc.13274
- [REF:worku2022] Worku DA (2022). SARS-CoV-2 associated immune dysregulation and COVID-associated pulmonary aspergillosis (CAPA): A cautionary tale. *Int. J. Mol. Sci.*, 23(6), 3228. doi:10.3390/ijms23063228
- [REF:pal2021] Pal R, Singh B, Bhadada SK, Banerjee M, Bhogal RS, Hage N & Kumar A (2021). COVID-19-associated mucormycosis: An updated systematic review of literature. *Mycoses*, 64(12), 1452–1459. doi:10.1111/myc.13338
- [REF:tsitsiklis2021] Tsitsiklis A, Zha BS, Byrne A, Devoe C, Levan S, Rackaityte E, ... & Langelier CR (2021). Impaired immune signaling and changes in the lung microbiome precede secondary bacterial pneumonia in COVID-19. *Res. Sq.* doi:10.21203/rs.3.rs-380803/v1
- Marrella V et al. (2024). Microbiota and immunity during respiratory infections: Lung and gut affair. *Int. J. Mol. Sci.* doi:10.3390/ijms25074051

- Jabeen MF et al. (2024). Species-level metagenomic and proteomic analysis of microbe-immune interactions in severe asthma. *Allergy*. doi:10.1111/all.16269
- Charalampous T et al. (2021). Respiratory metagenomics for secondary infection in COVID-19 ICUs. *Genome Med*. doi:10.1186/s13073-021-00991-y
- Woo S et al. (2020). Dynamics of Respiratory Microbiota during Mechanical Ventilation. *J. Clin. Med. Res*. doi:10.3390/jcm9030638
- Yeoh YK et al. (2021). Gut microbiota composition reflects disease severity in COVID-19. *Gut*. doi:10.1136/gutjnl-2020-323020
- Ren X et al. (2020). Intestinal microbiota in children with community-acquired pneumonia. *Microb. Pathog*. doi:10.1016/j.micpath.2020.104062

Ch. 8 – The Urogenital Microbiome

- [REF:ravel2011] Ravel J, Gajer P, Abdo Z, Schneider GM, Koenig SSK, McCulle SL, ... & Forney LJ (2011). Vaginal microbiome of reproductive-age women. *Proc. Natl. Acad. Sci. USA*, 108(Suppl 1), 4680–4687. doi:10.1073/pnas.1002611107
- [REF:gajer2012] Gajer P, Brotman RM, Bai G, Sakamoto J, Schütte UME, Zhong X, ... & Ravel J (2012). Temporal dynamics of the human vaginal microbiota. *Sci. Transl. Med.*, 4(132), 132ra52. doi:10.1126/scitranslmed.3003605
- [REF:smith2017] Smith SB & Ravel J (2017). The vaginal microbiota, host defence and reproductive physiology. *J. Physiol.*, 595(2), 451–463. doi:10.1113/JP271694
- [REF:chen2021] Chen X, Lu Y, Chen T & Li R (2021). The female vaginal microbiome in health and bacterial vaginosis. *Front. Cell. Infect. Microbiol.*, 11, 631972. doi:10.3389/fcimb.2021.631972
- [REF:levsagie2022] Lev-Sagie A, De Seta F, Verstraelen H, Ventolini G, Lonnee-Hoffmann R & Vieira-Baptista P (2022). The vaginal microbiome: II. Vaginal dysbiotic conditions. *J. Low. Genit. Tract Dis.*, 26(1), 79–84. doi:10.1097/LGT.0000000000000644
- [REF:price2020] Price TK, Hilt EE, Thomas-White K, Mueller ER, Wolfe AJ & Brubaker L (2020). The urobiome of continent adult women: a cross-sectional study. *BJOG*, 127(2), 193–201. doi:10.1111/1471-0528.15920
- [REF:wolfe2022] Joyce C, Halverson T, Gonzalez C, Brubaker L & Wolfe AJ (2022). The urobiomes of adult women with various lower urinary tract symptoms status differ: a re-analysis. *Front. Cell. Infect. Microbiol.*, 12, 860408. doi:10.3389/fcimb.2022.860408
- [REF:kenneally2022] Kenneally C, Murphy CP, Sleator RD & Culligan EP (2022). The urinary microbiome and biological therapeutics: novel therapies for urinary tract infections. *Microbiol. Res.*, 259, 127010. doi:10.1016/j.micres.2022.127010
- [REF:thomas-white2020] Thomas-White K, Taege S, Limeira R, Brincat C, Joyce C, Hilt EE, ... & Wolfe AJ (2020). Vaginal estrogen therapy is associated with increased *Lactobacillus* in the urine of postmenopausal women with overactive bladder symptoms. *Am. J. Obstet. Gynecol.*, 223(5), 727.e1–727.e11. doi:10.1016/j.ajog.2020.08.006
- [REF:gudnadottir2022] Gudnadottir U, Debelius JW, Du J, Hugerth LW, Danielsson H, Schuppe-Koistinen I, Fransson E & Brusselaers N (2022). The vaginal microbiome and the risk of preterm birth: a systematic review and network meta-analysis. *Sci. Rep.*, 12(1), 7926. doi:10.1038/s41598-022-12007-9
- [REF:huang2023] Huang C, Gin C, Fettweis J, Foxman B, Gelaye B, MacIntyre DA, ... & Callahan B (2023). Meta-analysis reveals the vaginal microbiome is a better predictor of earlier than later preterm birth. *BMC Biol.*, 21(1), 199. doi:10.1186/s12915-023-01702-2
- [REF:ihmp2019] The Integrative Human Microbiome Project Research Network Consor-

- tium (2019). The Integrative Human Microbiome Project. *Nature*, 569(7758), 641–648. doi:10.1038/s41586-019-1238-8
- [REF:plottel2011] Plottel CS & Blaser MJ (2011). Microbiome and malignancy. *Cell Host Microbe*, 10(4), 324–335. doi:10.1016/j.chom.2011.10.003
 - [REF:razavi2019] Razavi AC, Potts KS, Kelly TN & Bazzano LA (2019). Sex, gut microbiome, and cardiovascular disease risk. *Biol. Sex Differ.*, 10(1), 29. doi:10.1186/s13293-019-0240-z
 - [REF:jasarevic2016] Jašarević E, Morrison KE & Bale TL (2016). Sex differences in the gut microbiome-brain axis across the lifespan. *Philos. Trans. R. Soc. Lond. B Biol. Sci.*, 371(1688), 20150122. doi:10.1098/rstb.2015.0122
 - [REF:wu2024] Wu Z, Huang Y, Zhang R, Zheng C, You F, Wang M, Xiao C & Li X (2024). Sex differences in colorectal cancer: with a focus on sex hormone-gut microbiome axis. *Cell Commun. Signal.*, 22(1), 167. doi:10.1186/s12964-024-01549-2
 - Sadeghpour Heravi F (2024). Host-vaginal microbiota interaction: shaping the vaginal microenvironment. *Curr. Clin. Microbiol. Rep.* doi:10.1007/s40588-024-00227-8

Ch. 9 — The Oral Microbiome

- Samman M et al. (2022). Diet drinks and dental caries among U.S. adults. *Community Dent. Health*. doi:10.1922/CDH_00154Samman07

Ch. 10 — The Mycobiome

- [REF:iliev2012] Iliev ID et al. (2012). Interactions between commensal fungi and the C-type lectin receptor Dectin-1 influence colitis. *Science*, 336(6086), 1314–1317. doi:10.1126/science.1221789
- [REF:zhang2022] Zhang F, Aschenbrenner D, Yoo JY & Zuo T (2022). The gut mycobiome in health, disease, and clinical applications. *Lancet Microbe*, 3(12), e969–e983. doi:10.1016/S2666-5247(22)00203-8
- [REF:findley2013] Findley K, Oh J, Yang J, Conlan S et al. (2013). Topographic diversity of fungal and bacterial communities in human skin. *Nature*, 498(7454), 367–370. doi:10.1038/nature12171
- [REF:iliev2013] Iliev ID & Underhill DM (2013). Striking a balance: fungal commensalism versus pathogenesis. *Curr. Opin. Microbiol.*, 16(3), 366–373. doi:10.1016/j.mib.2013.05.004
- [REF:doron2021] Doron I, Leonardi I, Li XV et al. (2021). Human gut mycobiota tune immunity via CARD9-dependent induction of anti-fungal IgG antibodies. *Cell*, 184(4), 1017–1031.e14. doi:10.1016/j.cell.2021.01.016
- [REF:sokol2017] Sokol H, Leducq V, Aschard H et al. (2017). Fungal microbiota dysbiosis in IBD. *Gut*, 66(6), 1039–1048. doi:10.1136/gutjnl-2015-310746
- [REF:noverr2004] Noverr MC, Noggle RM, Toews GB & Huffnagle GB (2004). Role of antibiotics and fungal microbiota in driving pulmonary allergic responses. *Infect. Immun.*, 72(9), 4996–5003. doi:10.1128/IAI.72.9.4996-5003.2004
- [REF:noverr2005] Noverr MC, Falkowski NR, McDonald RA, McKenzie AN & Huffnagle GB (2005). Development of allergic airway disease in mice following antibiotic therapy and fungal microbiota increase. *Infect. Immun.*, 73(1), 30–38. doi:10.1128/IAI.73.1.30-38.2005
- [REF:jung2023] Jung WH (2023). Alteration in skin mycobiome due to atopic dermatitis and seborrheic dermatitis. *Biophys. Rev.*, 4(1), 011309. doi:10.1063/5.0136543
- [REF:belvoncikova2022] Belvončíková P et al. (2022). The human mycobiome: colonization, composition and the role in health and disease. *J. Fungi*, 8(10), 1046.

doi:10.3390/jof8101046

- [REF:wang2022] Wang Y-R, Zhu T, Kong F-Q et al. (2022). Infant mode of delivery shapes the skin mycobiome of prepubescent children. *Microbiol. Spectr.*, 10(5), e02267-22. doi:10.1128/spectrum.02267-22
- [REF:ward2017] Ward TL, Knights D & Gale CA (2017). Infant fungal communities: current knowledge and research opportunities. *BMC Med.*, 15(1), 30. doi:10.1186/s12916-017-0802-z
- [REF:gutierrez2022] Gutierrez MW, van Tilburg Bernardes E et al. (2022). “Molding” immunity — modulation of mucosal and systemic immunity by the intestinal mycobiome. *Mucosal Immunol.*, 15(4), 573–583. doi:10.1038/s41385-022-00515-w
- [REF:gutierrez2025] Gutierrez MW, van Tilburg Bernardes E, Ren E et al. (2025). Early-life gut mycobiome core species modulate metabolic health in mice. *Nat. Commun.*, 16(1), 1467. doi:10.1038/s41467-025-56743-8
- [REF:he2025] He S, Sun Y, Yu J et al. (2025). Dysbiotic gut fungi exacerbate *Klebsiella pneumoniae* lung infection via Dectin-1-mediated alveolar macrophage hyperactivation. *ISME J.*, 19(1). doi:10.1093/ismejo/wraf181
- [REF:hsu2023] Hsu C, Ghannoum M, Cominelli F & Di Martino L (2023). Mycobiome and inflammatory bowel disease: role in disease pathogenesis. *Inflamm. Bowel Dis.*, 29(3), 470–479. doi:10.1093/ibd/izac156
- D’Andreano S et al. (2021). Rapid identification of fungi with long amplicon nanopore sequencing. *Biol Methods Protoc.* doi:10.1093/biomethods/bpaa026
- Money NP (2021). Hyphal and mycelial consciousness: the concept of the fungal mind. *Fungal Biol.* doi:10.1016/j.funbio.2021.02.001
- Lovett B et al. (2020). Behavioral betrayal: How select fungal parasites enlist living insects. *PLoS Pathog.* doi:10.1371/journal.ppat.1008598

Ch. 11 — The Brain Microbiome

- Heller J (2024). The brain might also have a microbiome — what you need to know.
- Link CD (2021). Is there a brain microbiome? *Neurosci. Insights.* doi:10.1177/26331055211018709
- Yong SJ (2024). What If Something Else Lives in Your Brain? The Science and Secrets of Brain Microbiome.
- Hu Z et al. (2023). The remarkable complexity of the brain microbiome in health and disease. *bioRxiv.* doi:10.1101/2023.02.06.527297
- Moné Y et al. (2023). Evidence for a bacterial component in the etiology of Alzheimer’s disease. *Front. Cell. Infect. Microbiol.* doi:10.3389/fcimb.2023.1123228

Ch. 13 — The Gut-Brain Axis

- [REF:shatalina2024] Shatalina E, Whitehurst TS, Onwordi EC, Gilbert BJ, Rizzo G, Whittington A, Mansur A, Tsukada H, Reis Marques T, Natesan S, Rabiner EA, Wall MB & Howes OD (2024). Mitochondrial complex I density is associated with IQ and cognition in cognitively healthy adults: an in vivo [18F]BCPP-EF PET study. *EJNMMI Research*, 14, 41. doi:10.1186/s13550-024-01099-1
- [REF:geary2018] Geary DC (2018). Efficiency of mitochondrial functioning as the fundamental biological mechanism of general intelligence (g). *Psychological Review*, 125(6), 1028. doi:10.1037/rev0000124
- [REF:geary2019] Geary DC (2019). Mitochondria as the Linchpin of General Intelli-

- gence and the link between g, Health, and aging. *Journal of Intelligence*, 7(4), 25. doi:10.3390/jintelligence7040025
- [REF:picard2014] Picard M & McEwen BS (2014). Mitochondria impact brain function and cognition. *Proceedings of the National Academy of Sciences*, 111(1), 7–8. doi:10.1073/pnas.1321881111
 - [REF:tucker2018] Tucker D, Lu Y & Zhang Q (2018). From mitochondrial function to neuroprotection — an emerging role for methylene blue. *Molecular Neurobiology*, 55(6), 5137–5153. doi:10.1007/s12035-017-0712-2
 - [REF:terada2022] Terada T et al. (2022). Mitochondrial complex I abnormalities underlie neurodegeneration and cognitive decline in Alzheimer’s disease. *European Journal of Neurology*, 29(5), 1324–1334. doi:10.1111/ene.15266
 - Cox TO et al. (2026). Intestinal interoceptive dysfunction drives age-associated cognitive decline. *Nature*. doi:10.1038/s41586-026-10191-6
 - Cheng YT & Mazmanian SK (2026). Gut microbes affect cognition during ageing. doi:10.1038/d41586-026-00492-1
 - Cryan JF & Dinan TG (2012). Mind-altering microorganisms: the gut microbiota on brain and behaviour. *Nat. Rev. Neurosci.* doi:10.1038/nrn3346
 - Cryan JF et al. (2020). The gut microbiome in neurological disorders. *Lancet Neurol.* doi:10.1016/S1474-4422(19)30356-4
 - Torres-Fuentes C et al. (2017). The microbiota-gut-brain axis in obesity. *Lancet Gastroenterol. Hepatol.* doi:10.1016/S2468-1253(17)30147-4
 - Fekete M et al. (2024). Gut-brain axis modulation on cognitive health: prebiotics, probiotics, and symbiotics. *Nutrients*. doi:10.3390/nu16060789
 - Meyer K et al. (2022). Gut Microbiota and cognitive function in midlife. *JAMA Netw. Open*. doi:10.1001/jamanetworkopen.2021.43941
 - DOE/PNNL (2020). Scientists explore links between genetics, gut microbiome and memory. *Science Daily*.
 - Willyard C (2021). How gut microbes could drive brain disorders. *Nature*. doi:10.1038/d41586-021-00260-3
 - Mehta RS et al. (2022). Midlife antibiotic use and subsequent cognitive function in women. *PLoS One*. doi:10.1371/journal.pone.0264649

Ch. 14 — Behavioural Modulation

- Hakimi MA et al. (2017). Toxoplasma effectors targeting host signaling and transcription. *Clin. Microbiol. Rev.* doi:10.1128/CMR.00005-17
- Johnson HJ & Koshy AA (2020). Latent toxoplasmosis effects on rodents and humans. *MBio*. doi:10.1128/mBio.02164-19
- de Haan L et al. (2021). Toxoplasma gondii seropositivity and cognitive function: systematic review and meta-analysis. *JAMA Psychiatry*. doi:10.1001/jamapsychiatry.2021.1590
- Stock AK et al. (2017). Humans with latent toxoplasmosis display altered reward modulation. *Sci. Rep.* doi:10.1038/s41598-017-10926-6
- Ortiz-Guerrero G et al. (2020). Cognitive impairment and neurodegeneration by Toxoplasma gondii. *Brain Sci.* doi:10.3390/brainsci10060369
- Rougier S et al. (2017). Lifelong Persistence of Toxoplasma Cysts: A Questionable Dogma? *Trends Parasitol.* doi:10.1016/j.pt.2016.10.007
- Sutherland AL et al. (2015). Toxoplasma gondii in schizophrenia, bipolar disorder, and addiction: meta-analysis. *Acta Psychiatr. Scand.* doi:10.1111/acps.12423

Ch. 15 — Metabolism, Hormones, and the Endocrine Microbiome

- Yue F et al. (2022). Constant light exposure alters gut microbiota and aggravates PCOS-like traits. *Obesity*. doi:10.1002/oby.23380
- Sarmiento J et al. (2022). Nocturnal Light Pollution Induces Weight Gain and Reshapes Colonic Microbiota. *Int. J. Mol. Sci.* doi:10.3390/ijms23031673

Ch. 16–17 — Bacteriophages

- Carroll-Portillo A & Lin HC (2019). Bacteriophage and the innate immune system. *Microorganisms*. doi:10.3390/microorganisms7120625
- Dutilh BE et al. (2014). Highly abundant bacteriophage in unknown sequences of human faecal metagenomes. *Nat. Commun.* doi:10.1038/ncomms5498
- Bondy-Denomy J & Davidson AR (2014). When a virus is not a parasite: beneficial effects of prophages. *J. Microbiol.* doi:10.1007/s12275-014-4083-3
- Atanasova NS et al. (2015). Comparison of lipid-containing bacterial and archaeal viruses. *Adv. Virus Res.* doi:10.1016/bs.aivir.2014.11.005
- Zheludev IN et al. (2024). Viroid-like colonists of human microbiomes [“Obelisks”]. *bioRxiv*. doi:10.1101/2024.01.20.576352
- Sidik S (2024). ‘Wildly weird’ RNA bits discovered infesting the microbes in our guts. *Nature*. doi:10.1038/d41586-024-00266-7
- Thompson J (n.d.). Weird ‘obelisks’ found in human gut may be virus-like entities. *Sci. Am.*
- Raoult D & Forterre P (2008). Redefining viruses: lessons from Mimivirus. *Nat. Rev. Microbiol.* doi:10.1038/nrmicro1858

Ch. 18 — Phage Therapy

- Romero-Calle D et al. (2019). Bacteriophages as alternatives to antibiotics in clinical care. *Antibiotics*. doi:10.3390/antibiotics8030138
- Kushwaha SO et al. (2024). Bacteriophages as potential substitute for antibiotics: Comprehensive review. *Cell Biochem. Funct.* doi:10.1002/cbf.4022
- Petrenko VA (2024). Phage display’s prospects for early diagnosis of prostate cancer. *Viruses*. doi:10.3390/v16020277

Ch. 19 — Antibiotics: Collateral Damage

- Baldanzi G et al. (2026). Antibiotic use and gut microbiome composition: 14,979 individuals. *Nat. Med.* doi:10.1038/s41591-026-04284-y
- Muñoz KA et al. (2024). A Gram-negative-selective antibiotic that spares the gut microbiome. *Nature*. doi:10.1038/s41586-024-07502-0
- Althubyani AA et al. (2024). Antibiotic-induced neuropsychiatric toxicity. *Drugs Context*. doi:10.7573/dic.2024-3-3
- Livanos AE et al. (2016). Antibiotic-mediated gut microbiome perturbation accelerates T1D in mice. *Nat Microbiol.* doi:10.1038/nmicrobiol.2016.140
- Patrick DM et al. (2020). Decreasing antibiotic use, gut microbiota, and asthma incidence in children. *Lancet Respir. Med.* doi:10.1016/S2213-2600(20)30052-7
- Llor C et al. (2011). Impact of rapid antigen testing in acute pharyngitis on antibiotic prescription. *Br. J. Gen. Pract.* doi:10.3399/bjgp11X572436

- von Dach E et al. (2020). CRP-guided antibiotic duration in Gram-negative bacteremia. *JAMA*. doi:10.1001/jama.2020.6348
- Rhee C et al. (2020). Antibiotic-resistant pathogens in culture-proven sepsis. *JAMA Netw Open*. doi:10.1001/jamanetworkopen.2020.2899
- Ternák G et al. (2022). Antibiotic consumption patterns and diabetes prevalence in Europe. *Front. Endocrinol*. doi:10.3389/fendo.2022.870465
- Shuai M et al. (2022). Human Gut Antibiotic Resistome and Progression of Diabetes. *Adv. Sci*. doi:10.1002/advs.202104965
- Jabbar A & Sajjad-Ur-Rehman (2013). Antibiotic residues in meat, milk and eggs. *J Microb Biotech Food Sci*.
- Chowdhury S et al. (2015). Antibiotic residues in milk and eggs in Chittagong, Bangladesh. *Vet World*. doi:10.14202/vetworld.2015.467-471
- Boursi B et al. (2015). Recurrent antibiotic exposure may promote cancer formation. *Eur. J. Cancer*.
- Park SJ et al. (2021). Antibiotics use and diabetes incidence in Korean cohort. *Sci. Rep*. doi:10.1038/s41598-021-01125-5
- Nguyen K et al. (2013). Antibiotic use and resistance in Viet Nam. *BMC Public Health*. doi:10.1186/1471-2458-13-1158

Ch. 20 — Diet: Feeding the Ecosystem

Primary references cited in Chapter 20 draft

- Asnicar F et al. (2021). Microbiome connections with host metabolism and habitual diet from 1,098 deeply phenotyped individuals. *Nat. Med*. doi:10.1038/s41591-020-01183-8
- Bennet SMP et al. (2018). Multivariate modelling of faecal bacterial profiles and food intake in IBS. *Gut*. doi:10.1136/gutjnl-2016-313128
- Berry SE et al. (2020). Human postprandial responses to food and potential for precision nutrition. *Nat. Med*. doi:10.1038/s41591-020-0934-0
- Bielik V, Kolisek M (2021). Bioaccessibility and bioavailability of minerals in relation to a healthy gut microbiome. *Int. J. Mol. Sci*. doi:10.3390/ijms22136803
- Biesiekierski JR et al. (2013). No effects of gluten in patients with self-reported non-celiac gluten sensitivity after dietary reduction of FODMAPs. *Gastroenterology*. doi:10.1053/j.gastro.2013.04.051
- Chassaing B et al. (2015). Dietary emulsifiers impact the mouse gut microbiota promoting colitis and metabolic syndrome. *Nature*. doi:10.1038/nature14232
- Chassaing B et al. (2022). Randomized controlled-feeding study of dietary emulsifier CMC. *Gastroenterology*. doi:10.1053/j.gastro.2021.11.006
- Cortés-Martín A et al. (2020). Polyphenols and gut microbiota metabotypes. *Mol. Nutr. Food Res*. doi:10.1002/mnfr.201900952
- David LA et al. (2014). Diet rapidly and reproducibly alters the human gut microbiome. *Nature*. doi:10.1038/nature12820
- Davis DR, Epp MD, Riordan HD (2004). Changes in USDA food composition data for 43 garden crops, 1950 to 1999. *J. Am. Coll. Nutr*. doi:10.1080/07315724.2004.10719409
- Deehan EC et al. (2024). Effects of dietary fibre on metabolic health and obesity. *Nat. Rev. Gastroenterol. Hepatol*. doi:10.1038/s41575-023-00891-z
- Desai MS et al. (2016). A dietary fiber-deprived gut microbiota degrades the colonic mucus barrier. *Cell*. doi:10.1016/j.cell.2016.10.043

- Devkota S et al. (2012). Dietary-fat-induced taurocholic acid promotes pathobiont expansion and colitis. *Nature*. doi:10.1038/nature11225
- Furusawa Y et al. (2013). Commensal microbe-derived butyrate induces Treg differentiation. *Nature*. doi:10.1038/nature12721
- Ghosh TS et al. (2020). Mediterranean diet intervention alters the gut microbiome in older people (NU-AGE). *Gut*. doi:10.1136/gutjnl-2019-319654
- Halmos EP et al. (2014). Diets that differ in FODMAP content alter the colonic luminal microenvironment. *Gut*. doi:10.1136/gutjnl-2014-307264
- Koeth RA et al. (2013). Intestinal microbiota metabolism of L-carnitine promotes atherosclerosis. *Nat. Med.* doi:10.1038/nm.3145
- Koh A et al. (2016). From dietary fiber to host physiology: SCFAs as key bacterial metabolites. *Cell*. doi:10.1016/j.cell.2016.05.041
- Kolodziejczyk AA et al. (2019). Diet-microbiota interactions and personalized nutrition. *Nat. Rev. Microbiol.* doi:10.1038/s41579-019-0256-8
- Lebewohl B et al. (2017). Long-term gluten consumption and coronary heart disease risk. *BMJ*. doi:10.1136/bmj.j1892
- Lönnerdal B (2017). Excess iron intake as a factor in growth, infections, and development of infants and young children. *Am. J. Clin. Nutr.* doi:10.3945/ajcn.117.156042
- Mao Q et al. (2018). Ramazzini Institute 13-week pilot study on glyphosate and Roundup in Sprague-Dawley rats: effects on the microbiome. *Environ. Health*. doi:10.1186/s12940-018-0394-x
- Marco ML et al. (2017). Health benefits of fermented foods: microbiota and beyond. *Curr. Opin. Biotechnol.* doi:10.1016/j.copbio.2016.11.010
- McDonald D et al. (2018). American Gut: an open platform for citizen science microbiome research. *mSystems*. doi:10.1128/mSystems.00031-18
- Meslier V et al. (2020). Mediterranean diet intervention in overweight/obese subjects. *Gut*. doi:10.1136/gutjnl-2019-320438
- Mesnage R et al. (2021). Shotgun metagenomics and metabolomics of glyphosate or Roundup MON 52276 on gut microbiota of Sprague-Dawley rats. *Environ. Health Perspect.* doi:10.1289/EHP6990
- Mesnage R et al. (2022). Impacts of dietary exposure to pesticides on faecal microbiome metabolism in adult twins. *Environ. Health*. doi:10.1186/s12940-022-00860-0
- Motta EVS, Raymann K, Moran NA (2018). Glyphosate perturbs the gut microbiota of honey bees. *PNAS*. doi:10.1073/pnas.1803880115
- Paganini D et al. (2017). Prebiotic GOS mitigate adverse effects of iron fortification on the gut microbiome: RCT in Kenyan infants. *Gut*. doi:10.1136/gutjnl-2017-314418
- Paganini D et al. (2018). Iron-containing micronutrient powders modify the effect of oral antibiotics and increase post-antibiotic diarrhoea risk. *Gut*. doi:10.1136/gutjnl-2018-317399
- Palomar-Cros A et al. (2023). Dietary circadian rhythms and cardiovascular disease risk in NutriNet-Santé. *Nat. Commun.* doi:10.1038/s41467-023-43444-3
- Ranaivo H et al. (2022). Increasing the diversity of dietary fibers in daily-consumed bread modifies gut microbiota and metabolic profile. *Gut Microbes*. doi:10.1080/19490976.2022.2044722
- Reynolds A et al. (2019). Carbohydrate quality and human health: systematic reviews and meta-analyses. *Lancet*. doi:10.1016/S0140-6736(18)31809-9
- Schnorr SL et al. (2014). Gut microbiome of the Hadza hunter-gatherers. *Nat. Commun.* doi:10.1038/ncomms4654
- Smith PM et al. (2013). SCFAs regulate colonic Treg homeostasis. *Science*. doi:10.1126/science.1241165
- Smits SA et al. (2017). Seasonal cycling in the gut microbiome of the Hadza. *Science*.

doi:10.1126/science.aan4834

- Sonnenburg ED, Sonnenburg JL (2014). Starving our microbial self: consequences of MAC-deficient diet. *Cell Metab.* doi:10.1016/j.cmet.2014.07.003
- Sonnenburg ED et al. (2016). Diet-induced extinctions in the gut microbiota compound over generations. *Nature.* doi:10.1038/nature16504
- Staudacher HM, Whelan K (2017). The low FODMAP diet: mechanisms and efficacy in IBS. *Gut.* doi:10.1136/gutjnl-2017-313750
- Suez J et al. (2014). Artificial sweeteners induce glucose intolerance by altering gut microbiota. *Nature.* doi:10.1038/nature13793
- Suez J et al. (2022). Personalized microbiome-driven effects of non-nutritive sweeteners on human glucose tolerance. *Cell.* doi:10.1016/j.cell.2022.07.016
- Taylor BC et al. (2020). Consumption of fermented foods is associated with systematic differences in gut microbiome and metabolome. *mSystems.* doi:10.1128/mSystems.00901-19
- Thaïss CA et al. (2014). Transkingdom control of microbiota diurnal oscillations promotes metabolic homeostasis. *Cell.* doi:10.1016/j.cell.2014.09.048
- Thaïss CA et al. (2016). Microbiota diurnal rhythmicity programs host transcriptome oscillations. *Cell.* doi:10.1016/j.cell.2016.11.003
- Tomás-Barberán FA et al. (2017). Urolithins and metabolites. *Mol. Nutr. Food Res.* doi:10.1002/mnfr.201500901
- Turróni S et al. (2016). Fecal metabolome of the Hadza hunter-gatherers. *Sci. Rep.* doi:10.1038/srep32826
- Walsh L, Hill C, Ross RP (2023). Impact of glyphosate (Roundup) on the composition and functionality of the gut microbiome. *Gut Microbes.* doi:10.1080/19490976.2023.2263935
- Wang Z et al. (2011). Gut flora metabolism of phosphatidylcholine promotes cardiovascular disease. *Nature.* doi:10.1038/nature09922
- Wastyk HC et al. (2021). Gut-microbiota-targeted diets modulate human immune status. *Cell.* doi:10.1016/j.cell.2021.06.019
- Wu GD et al. (2011). Linking long-term dietary patterns with gut microbial enterotypes. *Science.* doi:10.1126/science.1208344
- Zeb F et al. (2020). Time-restricted feeding, metabolic risk, and circadian rhythm. *Br. J. Nutr.* doi:10.1017/S0007114519003428
- Zeevi D et al. (2015). Personalized nutrition by prediction of glycemic responses. *Cell.* doi:10.1016/j.cell.2015.11.001
- Zinöcker MK, Lindseth IA (2018). The Western diet-microbiome-host interaction and metabolic disease. *Nutrients.* doi:10.3390/nu10030365

Additional curated references (not yet cited in draft — reserve for possible inclusion)

- Li H et al. (2024). Resistant starch intake and weight loss by reshaping gut microbiota. *Nat Metab.* doi:10.1038/s42255-024-00988-y
- Lane MM et al. (2024). Ultra-processed food exposure and adverse health outcomes: umbrella review. *BMJ.* doi:10.1136/bmj-2023-077310
- Czarnecka K et al. (2021). Aspartame — True or False? *Nutrients.* doi:10.3390/nu13061957
- Mohr AE et al. (2024). Gut microbiome remodeling: protein pacing with intermittent fasting. *Nat. Commun.* doi:10.1038/s41467-024-48355-5
- Corbin KD et al. (2023). Host-diet-gut microbiome interactions and energy balance. *Nat. Commun.* doi:10.1038/s41467-023-38778-x
- Choi J et al. (2025). Yogurt promotes microbial diversity. *Front. Nutr.* doi:10.3389/fnut.2025.1609102

- Collins J et al. (2018). Dietary trehalose enhances virulence of epidemic *C. difficile*. *Nature*. doi:10.1038/nature25178
- Mesnage R & Antoniou MN (2017). Facts and fallacies in the debate on glyphosate toxicity. *Front. Public Health*. doi:10.3389/fpubh.2017.00316

Ch. 21 — Environment, Lifestyle, and the Modern World

- Huang D et al. (2022). Polystyrene microplastic exposure induces insulin resistance via dysbacteriosis. *Sci. Total Environ*. doi:10.1016/j.scitotenv.2022.155937
- Li D et al. (2020). Microplastic release from polypropylene feeding bottles. *Nature Food*. doi:10.1038/s43016-020-00171-y
- Zrimec J et al. (2021). Plastic-Degrading Potential across the Global Microbiome. *MBio*. doi:10.1128/mBio.02155-21
- Lu J et al. (2018). Triclosan induces multiple antibiotic resistance through genetic mutation. *Environ. Int*. doi:10.1016/j.envint.2018.06.004
- Short FL et al. (2021). Benzalkonium chloride antagonises aminoglycoside antibiotics. *EBio-Medicine*. doi:10.1016/j.ebiom.2021.103653
- Singh R et al. (2019). Lifestyle-Induced Microbial Gradients: An Indian Perspective. *Front. Microbiol*. doi:10.3389/fmicb.2019.02874
- Oduaran OH et al. (2020). Gut Microbiome of Rural and Urban South African Cohort. *bioRxiv*.
- Tomiyama AJ (2019). Stress and Obesity. *Annu. Rev. Psychol*. doi:10.1146/annurev-psych-010418-102936

Ch. 22 — Drugs Beyond Antibiotics

Primary references cited in Chapter 22 draft

Foundational cross-class screens

- [REF:maier2018] Maier L, Pruteanu M, Kuhn M, et al. (2018). Extensive impact of non-antibiotic drugs on human gut bacteria. *Nature*, 555(7698), 623–628. doi:10.1038/nature25979 (PMID: 29555994)
- [REF:vichvila2020] Vich Vila A, Collij V, Sanna S, et al. (2020). Impact of commonly used drugs on the composition and metabolic function of the gut microbiota. *Nat. Commun.*, 11(1), 362. doi:10.1038/s41467-019-14177-z (PMID: 31953381)
- [REF:garcia2024] Garcia-Santamarina S, et al. (2024). Emergence of community behaviors in the gut microbiota upon drug treatment. *Cell*, 187(22), 6346–6357.e20. doi:10.1016/j.cell.2024.08.037

PPIs

- [REF:imhann2016] Imhann F, Bonder MJ, Vich Vila A, et al. (2016). Proton pump inhibitors affect the gut microbiome. *Gut*, 65(5), 740–748. doi:10.1136/gutjnl-2015-310376 (PMID: 26657899)
- [REF:shin2002ppi] Shin JM, Sachs G. (2002). Restoration of acid secretion following treatment with proton pump inhibitors. *Gastroenterology*, 123(5), 1588–1597. doi:10.1053/gast.2002.36593 (PMID: 12404233)
- [REF:zhang2025ppi] Zhang M, et al. (2025). Proton pump inhibitor increases intestinal epithelial paracellular permeability via the p38-MAPK/NF-κB signaling pathway. *Mol. Biol*.

Rep., 53(1), 170. doi:10.1007/s11033-025-11336-y

- [REF:xu2025ppi] Xu J, et al. (2025). Unraveling gut microbiome interferences in cancer immunotherapy: a meta-analysis of diverse drug effects. *BMC Cancer*, 25, 1776. doi:10.1186/s12885-025-15094-9
- [REF:sharma2025ppi] Sharma S, Rawal P, Kalra S. (2025). Gut Guardianship: Recommendations for Rational Use of Proton Pump Inhibitors in Chronic Kidney Disease. *J. Pak. Med. Assoc.*, 75(11), 1822–1824. doi:10.47391/JPKMA.25-94

Metformin

- [REF:forslund2015] Forslund K, Hildebrand F, Nielsen T, et al. (2015). Disentangling type 2 diabetes and metformin treatment signatures in the human gut microbiota. *Nature*, 528(7581), 262–266. doi:10.1038/nature15766 (PMID: 26633628)
- [REF:su2025metformin] Su W, et al. (2025). The Microbial Regulation Spectrum of Metformin in Patients With Type 2 Diabetes: An Individual-Based Meta-Analysis of 1431 Participants. *J. Clin. Endocrinol. Metab.*, 110(8), 2383–2403. doi:10.1210/clinem/dgaf238
- [REF:attaye2024] Attaye I, et al. (2024). Oral Akkermansia muciniphila L2-7 augments glycemic control in type 2 diabetes. *iScience*, 27(8), 110455. doi:10.1016/j.isci.2024.110455
- [REF:xie2024metformin] Xie X, et al. (2024). Metformin reprograms tryptophan metabolism via gut microbiome-derived bile acid metabolites to ameliorate depression-like behaviors in mice. *Brain Behav. Immun.*, 123, 442–455. doi:10.1016/j.bbi.2024.09.014
- [REF:rosell2024] Rosell-Díaz M, et al. (2024). Metformin-induced changes in the gut microbiome and plasma metabolome are associated with cognition in men. *Metabolism*, 157, 155941. doi:10.1016/j.metabol.2024.155941

NSAIDs

- [REF:varela2025] Varela-Mendoza AA, et al. (2025). Kombucha Prevents Indomethacin-Induced Enteric Damage in Wistar Rat by Enhancing Epithelial Gut Barrier and Modulating Gut Microbiota. *Food Sci. Nutr.*, 13(8), e70804. doi:10.1002/fsn3.70804
- [REF:tuerxuntayi2025] Tuerxuntayi A, et al. (2025). Paeoniflorin attenuates ASA-induced enteropathy in simulated high-altitude hypoxia in rats by regulating intestinal inflammation, intestinal barrier, and intestinal flora. *Toxicol. Appl. Pharmacol.*, 506, 117641. doi:10.1016/j.taap.2025.117641

Statins

- [REF:oesterle2017] Oesterle A, Laufs U, Liao JK. (2017). Pleiotropic Effects of Statins on the Cardiovascular System. *Circ. Res.*, 120(3), 573–585. doi:10.1161/CIRCRESAHA.116.308537 (PMID: 28154098)
- [REF:cochrane2019eze] Defined SR, et al. (2019). Ezetimibe for the prevention of cardiovascular disease and all-cause mortality events. *Cochrane Database Syst. Rev.*, 2019(11), CD012502. doi:10.1002/14651858.CD012502.pub2
- [REF:khan2023pcsk9] Khan SU, et al. (2023). PCSK-9 Inhibitors and Cardiovascular Outcomes: A Systematic Review With Meta-Analysis. *Cardiovasc. Ther.*, 2023, 4653180. doi:10.1155/2023/4653180
- [REF:vieirasilva2020] Vieira-Silva S, Falony G, Belda E, et al. (2020). Statin therapy is associated with lower prevalence of gut microbiota dysbiosis. *Nature*, 581(7808), 310–315. doi:10.1038/s41586-020-2269-x (PMID: 32433607)
- [REF:verhaar2022] Verhaar BJH, Collard D, Prodan A, et al. (2022). Heterogeneity in statin responses explained by variation in the human gut microbiome. *Med*, 3(6), 388–405.e6.

doi:10.1016/j.medj.2022.05.002

- [REF:lind2024] Lind L, et al. (2024). Is LDL cholesterol associated with long-term mortality among primary prevention adults? A retrospective cohort study from a large healthcare system. *BMJ Open*, 14(3), e080151. doi:10.1136/bmjopen-2023-080151
- [REF:liu2021nhanes] Liu X, et al. (2021). Association between low density lipoprotein cholesterol and all-cause mortality: results from the NHANES 1999–2014. *Sci. Rep.*, 11(1), 22111. doi:10.1038/s41598-021-01738-w
- [REF:liu2025statin] Liu YH, et al. (2025). Gut microbiota profiles by LDL-C target achievement in statin-treated patients: A cross-sectional study. *Nutr. Metab. Cardiovasc. Dis.*, 36(3), 104465. doi:10.1016/j.numecd.2025.104465

Antipsychotics and antidepressants

- [REF:kamath2024] Kamath S, et al. (2024). The atypical antipsychotics lurasidone and olanzapine exert contrasting effects on the gut microbiome and metabolic function of rats. *Br. J. Pharmacol.*, 181(22), 4531–4545. doi:10.1111/bph.16507
- [REF:qian2023olanzapine] Qian L, et al. (2023). Longitudinal Gut Microbiota Dysbiosis Underlies Olanzapine-Induced Weight Gain. *Microbiol. Spectr.*, 11(4), e0005823. doi:10.1128/spectrum.00058-23
- [REF:zeng2024olanzapine] Zeng C, et al. (2024). B-GOS alleviates olanzapine-induced lipid disturbances in mice by enriching Akkermansia and upregulation of PGRMC1-Wnt signaling. *Food Chem. Toxicol.*, 185, 114490. doi:10.1016/j.fct.2024.114490
- [REF:bertossi2024] Bertossi F. (2024). A Possible Role of Akkermansia muciniphila in the Treatment of Olanzapine-Induced Weight Gain. *Cureus*, 16(3), e55733. doi:10.7759/cureus.55733
- [REF:dias2024] Dias MF, et al. (2024). Effects of antipsychotics on the gastrointestinal microbiota: A systematic review. *Psychiatry Res.*, 336, 115914. doi:10.1016/j.psychres.2024.115914

Laxatives

- [REF:tropini2018] Tropini C, Moss EL, Merrill BD, et al. (2018). Transient osmotic perturbation causes long-term alteration to the gut microbiota. *Cell*, 173(7), 1742–1754.e17. doi:10.1016/j.cell.2018.05.008 (PMID: 29906449)

Hormonal contraceptives and HRT

- [REF:hua2022] Hua X, et al. (2022). Longitudinal analysis of the impact of oral contraceptive use on the gut microbiome. *J. Med. Microbiol.*, 71(4). doi:10.1099/jmm.0.001512
- [REF:brito2025] Brito J, et al. (2025). Hormonal birth control is associated with altered gut microbiota β -diversity in physically active females across the menstrual cycle: a pilot trial. *J. Appl. Physiol.*, 138(3), 739–745. doi:10.1152/jappphysiol.00008.2025
- [REF:clapp2025] Clapp Organski A, et al. (2025). Oral Contraceptives Induce Time- and Intestinal Segment-Dependent Shifts in the Gut Microbiota. *Nutrients*, 17(16), 2591. doi:10.3390/nu17162591

Iron, vitamin D, fish oil, and herbal supplements

- [REF:baldi2024] Baldi A, et al. (2024). Effects of iron supplements and iron-containing micronutrient powders on the gut microbiome in Bangladeshi infants: a randomized controlled trial. *Nat. Commun.*, 15(1), 8640. doi:10.1038/s41467-024-53013-x
- [REF:tian2025iron] Tian C, et al. (2025). Intramuscular iron supplementation enhances intestinal barrier function in weaned piglets challenged with enterotoxigenic *E. coli*. *Front. Cell. Infect. Microbiol.*, 15, 1553639. doi:10.3389/fcimb.2025.1553639

- [REF:feng2025vitd] Feng X, et al. (2025). Vitamin D Deficiency Exacerbates d-Galactose-Induced Intestinal Barrier Dysfunction in a Murine Model of Intestinal Aging. *J. Agric. Food Chem.*, 73(39), 24725–24740. doi:10.1021/acs.jafc.5c07085
- [REF:wu2025vitd] Wu CW, et al. (2025). Vitamin D3 Improves Hypothalamic-Pituitary-Adrenal Axis Function, Immunological Responses, and Gut Dysbiosis in Sleep Desynchrony. *Brain Behav.*, 15(11), e71084. doi:10.1002/brb3.71084
- [REF:cao2025omega3] Cao B, et al. (2025). Interaction between dietary omega-3 polyunsaturated fatty acids, obesity and gut microbiota in preclinical models: A systematic review of randomized controlled trials. *Diabetes Obes. Metab.*, 27(9), 4643–4661. doi:10.1111/dom.16535
- [REF:roussel2025fishoil] Roussel C, et al. (2025). Human Gut Microbes Produce EPA- and DHA-Derived Oxylipins, but not N-Acyl-Ethanolamines, From Fish Oil. *FASEB J.*, 39(12), e70713. doi:10.1096/fj.202500752RR
- [REF:guo2023berberine] Guo HH, et al. (2023). Berberine is a potential alternative for metformin with good regulatory effect on lipids in treating metabolic diseases. *Biomed. Pharmacother.*, 163, 114754. doi:10.1016/j.biopha.2023.114754
- [REF:kong2020berberine] Kong WJ, et al. (2020). Berberine in the treatment of metabolism-related chronic diseases: A drug cloud (dCloud) effect to target multifactorial disorders. *Pharmacol. Ther.*, 209, 107496. doi:10.1016/j.pharmthera.2020.107496

Pharmacomicrobiomics — drugs metabolised by microbes

- [REF:zimmermann2019] Zimmermann M, Zimmermann-Kogadeeva M, Wegmann R, Goodman AL. (2019). Mapping human microbiome drug metabolism by gut bacteria and their genes. *Nature*, 570(7762), 462–467. doi:10.1038/s41586-019-1291-3 (PMID: 31158845)
- [REF:haiser2013] Haiser HJ, Gootenberg DB, Chatman K, Sirasani G, Balskus EP, Turnbaugh PJ. (2013). Predicting and manipulating cardiac drug inactivation by the human gut bacterium *Eggerthella lenta*. *Science*, 341(6143), 295–298. doi:10.1126/science.1235872 (PMID: 23869020)
- [REF:mainirekdal2019] Maini Rekdal V, Bess EN, Bisanz JE, Turnbaugh PJ, Balskus EP. (2019). Discovery and inhibition of an interspecies gut bacterial pathway for Levodopa metabolism. *Science*, 364(6445), eaau6323. doi:10.1126/science.aau6323 (PMID: 31196984)

Additional reserve references (not cited in current draft — kept for possible future inclusion)

- Hall RJ et al. (2024). Non-antibiotic pharmaceuticals are toxic against *E. coli*. *NPJ Antimicrob. Resist.* doi:10.1038/s44259-024-00028-5
- Chen H et al. (2025). Non-antibiotic medications and AMR in *E. coli*. *NPJ Antimicrob. Resist.* doi:10.1038/s44259-025-00144-w
- Imhann F et al. (2017). Proton pump inhibitors and other medications on gut microbiota. *Gut Microbes*. doi:10.1080/19490976.2017.1284732
- Le Bastard Q et al. (2021). Impact of non-antibiotic drugs on intestinal microbiome. *Expert Rev. Mol. Diagn.* doi:10.1080/14737159.2021.1952075
- Pyo JH et al. (2021). Proton pump inhibitors and fatty liver disease risk. *J. Gastroenterol. Hepatol.* doi:10.1111/jgh.15236
- Wu H et al. (2017). Metformin alters the gut microbiome in treatment-naïve T2D. *Nat. Med.* doi:10.1038/nm.4345

Ch. 23 — Probiotics: Promise, Hype, and Hard Truths

- Suez J et al. (2018). Post-antibiotic gut mucosal microbiome reconstitution is impaired by probiotics. *Cell*. doi:10.1016/j.cell.2018.08.047
- Butler CC et al. (2020). Probiotic Use and Antibiotic Administration in Care Home Residents. *JAMA*. doi:10.1001/jama.2020.8556
- Li G et al. (2023). Probiotics supplementation and glycaemic control in T2DM: meta-analysis. *J. Transl. Med.* doi:10.1186/s12967-023-04306-0
- Gao K et al. (2023). Bifidobacterium longum WHH2270 ameliorates type 2 diabetes in rats. *J. Food Sci.* doi:10.1111/1750-3841.16727
- Forsyth CB et al. (2009). Lactobacillus GG treatment ameliorates alcohol-induced liver injury. *Alcohol*. doi:10.1016/j.alcohol.2008.12.009
- Maguire M & Maguire G (2019). Gut dysbiosis, leaky gut, and intestinal epithelial proliferation in neurological disorders. *Rev. Neurosci.* doi:10.1515/revneuro-2018-0024

Ch. 24 — Faecal Microbiota Transplantation

- Udayappan SD et al. (2014). Intestinal microbiota and FMT as treatment for insulin resistance and T2DM. *Clin. Exp. Immunol.* doi:10.1111/cei.12293
- [REF:eiseman1958] Eiseman B, Silen W, Bascom GS & Kauvar AJ (1958). Fecal enema as an adjunct in the treatment of pseudomembranous enterocolitis. *Surgery*, 44(5), 854–859.
- [REF:gough2011] Gough E, Shaikh H & Manges AR (2011). Systematic review of intestinal microbiota transplantation (fecal bacteriotherapy) for recurrent *Clostridium difficile* infection. *Clin. Infect. Dis.*, 53(10), 994–1002. doi:10.1093/cid/cir632
- [REF:vannood2013] van Nood E, Vrieze A, Nieuwdorp M et al. (2013). Duodenal infusion of donor feces for recurrent *Clostridium difficile*. *N. Engl. J. Med.*, 368(5), 407–415. doi:10.1056/NEJMoa1205037
- [REF:vrieze2012] Vrieze A, Van Nood E, Holleman F et al. (2012). Transfer of intestinal microbiota from lean donors increases insulin sensitivity in individuals with metabolic syndrome. *Gastroenterology*, 143(4), 913–916.e7. doi:10.1053/j.gastro.2012.06.031
- [REF:alang2015] Alang N & Kelly CR (2015). Weight gain after fecal microbiota transplantation. *Open Forum Infect. Dis.*, 2(1), ofv004. doi:10.1093/ofid/ofv004
- [REF:paramsothy2017] Paramsothy S, Kamm MA, Kaakoush NO et al. (2017). Multidonor intensive faecal microbiota transplantation for active ulcerative colitis: a randomised placebo-controlled trial. *Lancet*, 389(10075), 1218–1228. doi:10.1016/S0140-6736(17)30182-4
- [REF:cammarota2017] Cammarota G, Ianiro G, Tilg H et al. (2017). European consensus conference on faecal microbiota transplantation in clinical practice. *Gut*, 66(4), 569–580. doi:10.1136/gutjnl-2016-313017
- [REF:kao2017] Kao D, Roach B, Silva M et al. (2017). Effect of oral capsule- vs colonoscopy-delivered fecal microbiota transplantation on recurrent *Clostridium difficile* infection: a randomized clinical trial. *JAMA*, 318(20), 1985–1993. doi:10.1001/jama.2017.17077
- [REF:kang2017] Kang D-W, Adams JB, Gregory AC et al. (2017). Microbiota Transfer Therapy alters gut ecosystem and improves gastrointestinal and autism symptoms: an open-label study. *Microbiome*, 5(1), 10. doi:10.1186/s40168-016-0225-7
- [REF:kang2019] Kang D-W, Adams JB, Coleman DM et al. (2019). Long-term benefit of Microbiota Transfer Therapy on autism symptoms and gut microbiota. *Sci. Rep.*, 9(1), 5821. doi:10.1038/s41598-019-42183-0

- [REF:defilipp2019] DeFilipp Z, Bloom PP, Torres Soto M et al. (2019). Drug-resistant *E. coli* bacteremia transmitted by fecal microbiota transplant. *N. Engl. J. Med.*, 381(21), 2043–2050. doi:10.1056/NEJMoa1910437
- [REF:ekekezie2020] Ekekezie C, Perler BK, Wexler A, Duff C, Lillis CJ & Kelly CR (2020). Understanding the scope of do-it-yourself fecal microbiota transplant. *Am. J. Gastroenterol.*, 115(4), 603–607. doi:10.14309/ajg.0000000000000499
- [REF:ianiro2020covid] Ianiro G, Mullish BH, Kelly CR et al. (2020). Reorganisation of faecal microbiota transplant services during the COVID-19 pandemic. *Gut*, 69(9), 1555–1563. doi:10.1136/gutjnl-2020-321829
- [REF:ianiro2020] Ianiro G, Bibbò S, Porcari S et al. (2021). Fecal microbiota transplantation for recurrent *C. difficile* infection in patients with inflammatory bowel disease: experience of a large-volume European FMT center. *Gut Microbes*, 13(1), 1994832. doi:10.1080/19490976.2021.1994832
- [REF:davar2021] Davar D, Dzutsev AK, McCulloch JA et al. (2021). Fecal microbiota transplant overcomes resistance to anti-PD-1 therapy in melanoma patients. *Science*, 371(6529), 595–602. doi:10.1126/science.abf3363
- [REF:degroot2020] de Groot P, Nikolic T, Pellegrini S et al. (2021). Faecal microbiota transplantation halts progression of human new-onset type 1 diabetes in a randomised controlled trial. *Gut*, 70(1), 92–105. doi:10.1136/gutjnl-2020-322630
- [REF:feuerstadt2022] Feuerstadt P, Louie TJ, Lashner B et al. (2022). SER-109, an oral microbiome therapy for recurrent *Clostridioides difficile* infection. *N. Engl. J. Med.*, 386(3), 220–229. doi:10.1056/NEJMoa2106516
- [REF:khanna2025] Khanna S, Yoho D, Van Handel D et al. (2025). Prevention of recurrent *C. difficile* infection by fecal microbiota, live-jslm (REBYOTA) administered via colonoscopy: 6-month data from the single-arm phase IIIb CDI-SCOPE trial. *Ther. Adv. Gastroenterol.*, 18, 17562848251396744. doi:10.1177/17562848251396744

Ch. 25 — Precision Microbiome Engineering

- [REF:depommier2019] Depommier C, Everard A, Druart C et al. (2019). Supplementation with *Akkermansia muciniphila* in overweight and obese human volunteers: a proof-of-concept exploratory study. *Nat. Med.*, 25(7), 1096–1103. doi:10.1038/s41591-019-0495-2
- [REF:kurtz2019] Kurtz CB, Millet YA, Puurunen MK et al. (2019). An engineered *E. coli* Nissle improves hyperammonemia and survival in mice and shows dose-dependent exposure in healthy humans. *Sci. Transl. Med.*, 11(475), eaau7975. doi:10.1126/scitranslmed.aau7975
- [REF:louie2023] Louie T, Golan Y, Khanna S et al. (2023). VE303, a defined bacterial consortium, for prevention of recurrent *Clostridioides difficile* infection: a randomized clinical trial. *JAMA*, 329(16), 1356–1366. doi:10.1001/jama.2023.4314
- [REF:synlogic2023] Synlogic Inc. (2023). Synlogic announces discontinuation of Synpheny-3 Phase 3 clinical trial of labafenogene marselecobac (SYNB1934) for the treatment of phenylketonuria. Company press release, 18 August 2023.
- [REF:locus2024] Kim P, Sanborn M, Nixon J et al. (2024). A randomized, double-blind, placebo-controlled Phase 1b/2a study of LBP-EC01, a CRISPR-Cas3-enhanced bacteriophage cocktail, for the suppression of *E. coli* in the urinary tract of patients with recurrent urinary tract infections (ELIMINATE trial). Company-reported data; ClinicalTrials.gov NCT05488340.
- [REF:eligo2025] Eligo Bioscience (2025). Pipeline and investor update: EBO05 programme advancing towards first-in-human studies. Company release.

- [REF:murray2022] Murray CJL, Ikuta KS, Sharara F et al. (2022). Global burden of bacterial antimicrobial resistance in 2019: a systematic analysis. *Lancet*, 399(10325), 629–655. doi:10.1016/S0140-6736(21)02724-0
- [REF:lagier2016] Lagier J-C, Khelaifia S, Alou MT et al. (2016). Culture of previously uncultured members of the human gut microbiota by culturomics. *Nat. Microbiol.*, 1, 16203. doi:10.1038/nmicrobiol.2016.203
- [REF:bikard2014] Bikard D, Euler CW, Jiang W et al. (2014). Exploiting CRISPR-Cas nucleases to produce sequence-specific antimicrobials. *Nat. Biotechnol.*, 32(11), 1146–1150. doi:10.1038/nbt.3043
- [REF:citorik2014] Citorik RJ, Mimee M & Lu TK (2014). Sequence-specific antimicrobials using efficiently delivered RNA-guided nucleases. *Nat. Biotechnol.*, 32(11), 1141–1145. doi:10.1038/nbt.3011
- [REF:neil2021] Neil K, Allard N, Roy P et al. (2021). High-efficiency delivery of CRISPR-Cas9 by engineered probiotics enables precise microbiome editing. *Mol. Syst. Biol.*, 17(10), e10335. doi:10.15252/msb.202110335
- [REF:pdaia2023] U.S. Food and Drug Administration, CBER (2016, updated 2023). Early clinical trials with live biotherapeutic products: chemistry, manufacturing, and control information — guidance for industry.
- [REF:hanson2021] Hanson BM & Weinstock GM (2021). The importance of the microbiome in epidemiologic research. *Ann. Epidemiol.*, 35, 75–79. doi:10.1016/j.annepidem.2019.05.007

Ch. 26 — Inflammatory and Autoimmune Disease

Core references used in Chapter 26 draft

- [REF:iliev2025] Iliev ID, Ananthakrishnan AN & Guo CJ (2025). Microbiota in inflammatory bowel disease: mechanisms of disease and therapeutic opportunities. *Nat. Rev. Microbiol.*, 23(8), 509–524. doi:10.1038/s41579-025-01163-0
- [REF:gevers2014] Gevers D, Kugathasan S, Denson LA et al. (2014). The treatment-naive microbiome in new-onset Crohn’s disease. *Cell Host Microbe*, 15(3), 382–392. doi:10.1016/j.chom.2014.02.005
- [REF:haneishi2023] Haneishi Y, Furuya Y, Hasegawa M et al. (2023). Inflammatory bowel diseases and gut microbiota. *Int. J. Mol. Sci.*, 24(4), 3817. doi:10.3390/ijms24043817
- [REF:darfeuille2004] Darfeuille-Michaud A, Boudeau J, Bulois P et al. (2004). High prevalence of adherent-invasive *Escherichia coli* associated with ileal mucosa in Crohn’s disease. *Gastroenterology*, 127(2), 412–421. doi:10.1053/j.gastro.2004.04.061
- [REF:henke2019] Henke MT, Kenny DJ, Cassilly CD, Vlamakis H, Xavier RJ & Clardy J (2019). *Ruminococcus gnavus*, a member of the human gut microbiome associated with Crohn’s disease, produces an inflammatory polysaccharide. *Proc. Natl. Acad. Sci. USA*, 116(26), 12672–12677. doi:10.1073/pnas.1904099116
- [REF:hugot2001] Hugot JP, Chamaillard M, Zouali H et al. (2001). Association of NOD2 leucine-rich repeat variants with susceptibility to Crohn’s disease. *Nature*, 411(6837), 599–603. doi:10.1038/35079107
- [REF:leibovitzh2024] Leibovitzh H, Nayeri S, Borowski K et al. (2024). Inflammatory bowel disease associated with primary sclerosing cholangitis is associated with an altered gut microbiome and bile acid profile. *J. Crohns Colitis*, 18(12), 1957–1966. doi:10.1093/ecco-jcc/jjae096
- [REF:sendid2023] Sendid B, Cornu M, Cordier C, Bouckaert J, Colombel JF & Poulain D

- (2023). From ASCA breakthrough in Crohn's disease and *Candida albicans* research to thirty years of investigations about their meaning in human health. *Autoimmun. Rev.*, 23(2), 103486. doi:10.1016/j.autrev.2023.103486
- [REF:fuchtbauer2025] Fuchtbauer JD, Berentsen EM, Lindhardt CL et al. (2025). Gut mycobioime in inflammatory bowel disease: a systematic review of case-control studies. *Inflamm. Bowel Dis.*, 31(11), 3213–3224. doi:10.1093/ibd/izaf217
 - [REF:hsu2023] Hsu C, Ghannoum M, Cominelli F & Di Martino L (2023). Mycobioime and inflammatory bowel disease: role in disease pathogenesis, current approaches and novel nutritional-based therapies. *Inflamm. Bowel Dis.*, 29(3), 470–479. doi:10.1093/ibd/izac156
 - [REF:shao2022] Shao TY, Haslam DB, Bennett RJ & Way SS (2022). Friendly fungi: symbiosis with commensal *Candida albicans*. *Trends Immunol.*, 43(9), 706–717. doi:10.1016/j.it.2022.07.003
 - [REF:fasano2012] Fasano A (2012). Leaky gut and autoimmune diseases. *Clin. Rev. Allergy Immunol.*, 42(1), 71–78. doi:10.1007/s12016-011-8291-x
 - [REF:divincenzo2023] Di Vincenzo F, Del Gaudio A, Petito V, Lopetuso LR & Scaldaferri F (2023). Gut microbiota, intestinal permeability, and systemic inflammation: a narrative review. *Intern. Emerg. Med.*, 19(2), 275–293. doi:10.1007/s11739-023-03374-w
 - [REF:mu2017] Mu Q, Kirby J, Reilly CM & Luo XM (2017). Leaky gut as a danger signal for autoimmune diseases. *Front. Immunol.*, 8, 598. doi:10.3389/fimmu.2017.00598
 - [REF:vatanen2018] Vatanen T, Franzosa EA, Schwager R et al. (2018). The human gut microbiome in early-onset type 1 diabetes from the TEDDY study. *Nature*, 562(7728), 589–594. doi:10.1038/s41586-018-0620-2
 - [REF:stewart2018teddy] Stewart CJ, Ajami NJ, O'Brien JL et al. (2018). Temporal development of the gut microbiome in early childhood from the TEDDY study. *Nature*, 562(7728), 583–588. doi:10.1038/s41586-018-0617-x
 - [REF:luppi2024] Luppi S, Aldegheri L, Azzalini E et al. (2024). Unravelling the role of gut and oral microbiota in the pediatric population with type 1 diabetes mellitus. *Int. J. Mol. Sci.*, 25(19), 10611. doi:10.3390/ijms251910611
 - [REF:asantebaadu2025] Asante Baadu F, Yalley C, Ohene Yeboah AR et al. (2025). Microbiome imbalance and pediatric type 1 diabetes mellitus: an updated systematic review of gut dysbiosis evidence. *Cureus*, 17(8), e89279. doi:10.7759/cureus.89279
 - [REF:jangi2016] Jangi S, Gandhi R, Cox LM et al. (2016). Alterations of the human gut microbiome in multiple sclerosis. *Nat. Commun.*, 7, 12015. doi:10.1038/ncomms12015
 - [REF:ordonezrodriguez2023] Ordoñez-Rodríguez A, Roman P, Rueda-Ruzafa L, Campos-Rios A & Cardona D (2023). Changes in gut microbiota and multiple sclerosis: a systematic review. *Int. J. Environ. Res. Public Health*, 20(5), 4624. doi:10.3390/ijerph20054624
 - [REF:schwerdtfeger2025] Schwerdtfeger LA, Jangi S, Yeh WZ et al. (2025). Gut microbiota and metabolites are linked to disease progression in multiple sclerosis. *Cell Rep. Med.*, 6(4), 102055. doi:10.1016/j.xcrm.2025.102055
 - [REF:scher2013] Scher JU, Sczesnak A, Longman RS et al. (2013). Expansion of intestinal *Prevotella copri* correlates with enhanced susceptibility to arthritis. *eLife*, 2, e01202. doi:10.7554/eLife.01202
 - [REF:jiang2022] Jiang L, Shang M, Yu S et al. (2022). A high-fiber diet synergizes with *Prevotella copri* and exacerbates rheumatoid arthritis. *Cell. Mol. Immunol.*, 19(12), 1414–1424. doi:10.1038/s41423-022-00934-6
 - [REF:abdelsalam2023] Abdelsalam NA, Hegazy SM & Aziz RK (2023). The curious case of *Prevotella copri*. *Gut Microbes*, 15(2), 2249152. doi:10.1080/19490976.2023.2249152
 - [REF:kim2023] Kim CH (2023). Complex regulatory effects of gut microbial short-chain

fatty acids on immune tolerance and autoimmunity. *Cell. Mol. Immunol.*, 20(4), 341–350. doi:10.1038/s41423-023-00987-1

- [REF:catassi2024] Catassi G, Aloï M, Giorgio V, Gasbarrini A, Cammarota G & Ianiro G (2024). The role of microbiome in the development of gluten-related disorders. *Best Pract. Res. Clin. Gastroenterol.*, 72, 101951. doi:10.1016/j.bpg.2024.101951
- [REF:elmouzan2022] El Mouzan M, Assiri A, Al Sarkhy A et al. (2022). Microbiota profile of new-onset celiac disease in children in Saudi Arabia. *Gut Pathog.*, 14(1), 37. doi:10.1186/s13099-022-00493-1
- [REF:belvoncikova2022] Belvoncikova P, Maronek M & Gardlik R (2022). Gut dysbiosis and fecal microbiota transplantation in autoimmune diseases. *Int. J. Mol. Sci.*, 23(18), 10729. doi:10.3390/ijms231810729
- [REF:aleman2023] Aleman RS, Moncada M & Aryana KJ (2023). Leaky gut and the ingredients that help treat it: a review. *Molecules*, 28(2), 619. doi:10.3390/molecules28020619

Additional references (not cited in current draft; kept for later sub-sections on SLE, ankylosing spondylitis, and the STI-antibiotic-autoimmunity tangent)

- Tomofuji Y et al. (2021). Gut microbiome of systemic lupus erythematosus in Japanese. *Ann. Rheum. Dis.* doi:10.1136/annrheumdis-2021-220687
- Huang R et al. (2020). Intestinal microbiome in ankylosing spondylitis patients. *J. Med. Microbiol.* doi:10.1099/jmm.0.001107
- Dedrick S et al. (2020). Gut microbiota and environmental factors in T1D pathogenesis. *Front. Endocrinol.* doi:10.3389/fendo.2020.00078
- Van Dijck C et al. (2021). Antibacterial mouthwash for STIs in MSM on PrEP (PREGo). *Lancet Infect. Dis.* doi:10.1016/S1473-3099(20)30778-7
- Kenyon C et al. (2023). Asymptomatic STI management in Europe. *Lancet Reg. Health Eur.* doi:10.1016/j.lanepe.2023.100743

Ch. 27 — Metabolic Disease and Obesity

Note: This is the primary home for the diabetes/obesity references. The bibliography is rich here — reflecting the author’s research interest — and will provide substantial material for a detailed chapter.

- Qin J et al. (2012). Metagenome-wide association study of gut microbiota in T2D. *Nature*. doi:10.1038/nature11450
- de Vos WM & Nieuwdorp M (2013). A gut prediction. *Nature*. doi:10.1038/nature12251
- Hartstra AV et al. (2014). Role of the Microbiome in Obesity and T2D. *Diabetes Care*. doi:10.2337/dc14-0769
- Tilg H & Moschen AR (2014). Microbiota and diabetes: an evolving relationship. *Gut*. doi:10.1136/gutjnl-2014-306928
- Barlow GM et al. (2015). Role of the Gut Microbiome in Obesity and DM. *Nutr. Clin. Pract.* doi:10.1177/0884533615609896
- Cai L et al. (2015). T2D Biomarkers of Human Gut Microbiota. *PLoS One*. doi:10.1371/journal.pone.0140827
- Sabatino A et al. (2017). Intestinal Microbiota in T2D and CKD. *Curr. Diab. Rep.* doi:10.1007/s11892-017-0841-z
- Tong X et al. (2018). Gut Microbiota during amelioration of T2D with hyperlipidemia by metformin. *MBio*. doi:10.1128/mBio.02392-17

- Liu G et al. (2018). Pumpkin polysaccharide modifies gut microbiota during alleviation of T2D. *Int. J. Biol. Macromol.* doi:10.1016/j.ijbiomac.2018.04.127
- Nie X et al. (2020). Gut microbiome and visceral fat accumulation. *Comput. Struct. Biotechnol. J.* doi:10.1016/j.csbj.2020.09.026
- Liu X et al. (2021). GWAS for gut metagenome in Chinese adults. *Cell Discov.* doi:10.1038/s41421-020-00239-w
- Wang Z et al. (2024). Gut Microbiota and blood metabolites related to fiber and T2D. *Circ. Res.* doi:10.1161/circresaha.123.323634
- Peng D et al. (2023). Dendrobium officinale derivatives in a T2D mouse model. *Food Chem.* doi:10.1016/j.foodchem.2022.133974
- Neri-Rosario D et al. (2023). Dysbiosis signatures and T2D progression: machine learning. *Front. Endocrinol.* doi:10.3389/fendo.2023.1170459
- Gravdal K et al. (2023). Gut microbiota in pre-diabetes and treatment-naïve T2D. *BMC Endocrin. Disord.* doi:10.1186/s12902-023-01432-0
- Hu J et al. (2023). Gut microbiota signatures in T1D and T2D. *EClinicalMedicine.* doi:10.1016/j.eclinm.2023.102132
- Zhang L et al. (2023). Intestinal Microbiota in Early and Late Stages of DKD. *Microbiol Spectr.* doi:10.1128/spectrum.00382-23
- Lu X et al. (2023). Gut microbiota in diabetic nephropathy. *Sci. Rep.* doi:10.1038/s41598-023-39444-4
- Chen X et al. (2023). Dendrobium officinale polysaccharide alleviates T2DM. *J. Agric. Food Chem.* doi:10.1021/acs.jafc.3c02429
- Cuadrat RRC et al. (2023). Gut microbiota and vascular stiffness. *Sci. Rep.* doi:10.1038/s41598-023-40178-6
- Antwi J (2023). Precision Nutrition for Obesity and T2D. *Curr. Nutr. Rep.* doi:10.1007/s13668-023-00491-y
- Ofosu FK et al. (2023). Fermented sorghum improves T2D via gut microbiota. *J. Funct. Foods.* doi:10.1016/j.jff.2023.105666
- Han CY et al. (2023). Benaglutide effects on Gut Microbiota in T2DM. *Diab. Metab. Syndr. Obes.* doi:10.2147/DMSO.S418757
- Niu X et al. (2023). Metformin + liraglutide and gut microbiota in T2D. *Int. Microbiol.* doi:10.1007/s10123-023-00380-y
- Zhou H et al. (2023). Diet composition influences cholesterol via gut Microbiota in ApoE^{-/-} mice. *J. Agric. Food Chem.* doi:10.1021/acs.jafc.3c00810
- Pierce DR et al. (2023). Experimental hookworm infection and insulin resistance. *Nat. Commun.* doi:10.1038/s41467-023-40263-4
- Hughes S (2022). Gut microbiome product implicated in heart failure risk.
- Melville NA (2022). Gut microbiome species predict T2D.
- Jie Z et al. (2023). Three-bacteria consortium inhibits atherosclerosis in mice. *iScience.* doi:10.1016/j.isci.2023.106960

References added during Ch. 27 drafting (new citations)

- Ley RE, Bäckhed F, Turnbaugh P, et al. (2005). Obesity alters gut microbial ecology. *Proc Natl Acad Sci USA.* 102(31):11070–5. doi:10.1073/pnas.0504978102
- Ley RE, Turnbaugh PJ, Klein S, Gordon JI (2006). Human gut microbes associated with obesity. *Nature.* 444(7122):1022–3. doi:10.1038/4441022a
- Turnbaugh PJ, Ley RE, Mahowald MA, et al. (2006). An obesity-associated gut mi-

- crobiome with increased capacity for energy harvest. *Nature*. 444(7122):1027–31. doi:10.1038/nature05414
- Le Chatelier E, Nielsen T, Qin J, et al. (2013). Richness of human gut microbiome correlates with metabolic markers. *Nature*. 500(7464):541–6. doi:10.1038/nature12506
 - Vrieze A, Van Nood E, Holleman F, et al. (2012). Transfer of intestinal microbiota from lean donors increases insulin sensitivity in individuals with metabolic syndrome. *Gastroenterology*. 143(4):913–916.e7. doi:10.1053/j.gastro.2012.06.031
 - Forslund K, Hildebrand F, Nielsen T, et al. (2015). Disentangling type 2 diabetes and metformin treatment signatures in the human gut microbiota. *Nature*. 528(7581):262–266. doi:10.1038/nature15766
 - Depommier C, Everard A, Druart C, et al. (2019). Supplementation with *Akkermansia muciniphila* in overweight and obese human volunteers: a proof-of-concept exploratory study. *Nat Med*. 25(7):1096–1103. doi:10.1038/s41591-019-0495-2
 - Witkowski M, Weeks TL, Hazen SL (2020). Gut Microbiota and Cardiovascular Disease. *Circ Res*. 127(4):553–570. doi:10.1161/CIRCRESAHA.120.316242
 - Zhu T, Goodarzi MO (2020). Metabolites Linking the Gut Microbiome with Risk for Type 2 Diabetes. *Curr Nutr Rep*. 9(2):83–93. doi:10.1007/s13668-020-00307-3
 - Münzker J, Haase N, Till A, et al. (2022). Functional changes of the gastric bypass microbiota reactivate thermogenic adipose tissue and systemic glucose control via intestinal FXR-TGR5 crosstalk in diet-induced obesity. *Microbiome*. 10(1):96. doi:10.1186/s40168-022-01264-5
 - Canyelles M, Borràs C, Rotllan N, et al. (2023). Gut Microbiota-Derived TMAO: A Causal Factor Promoting Atherosclerotic Cardiovascular Disease? *Int J Mol Sci*. 24(3):1940. doi:10.3390/ijms24031940
 - Attaye I, et al. (2024). Oral *Akkermansia muciniphila* L2-7 augments glycemic control in type 2 diabetes. *iScience*. 27(8):110455. doi:10.1016/j.isci.2024.110455
 - da Silva RS, de Paiva IHR, Mendonça IP, et al. (2024). Anorexigenic and anti-inflammatory signaling pathways of semaglutide via the microbiota-gut-brain axis in obese mice. *Inflammopharmacology*. 33(2):845–864. doi:10.1007/s10787-024-01603-y
 - Benedé-Ubieto R, Cubero FJ, Nevzorova YA (2024). Breaking the barriers: the role of gut homeostasis in Metabolic-Associated Steatotic Liver Disease (MASLD). *Gut Microbes*. 16(1):2331460. doi:10.1080/19490976.2024.2331460
 - Long Q, Luo F, Li B, et al. (2024). Gut microbiota and metabolic biomarkers in metabolic dysfunction-associated steatotic liver disease. *Hepatol Commun*. 8(3). doi:10.1097/HC9.0000000000000310
 - Gofron KK, Wasilewski A, Małgorzewicz S (2025). Effects of GLP-1 Analogues and Agonists on the Gut Microbiota: A Systematic Review. *Nutrients*. 17(8):1303. doi:10.3390/nu17081303
 - Zhao S, Lin H, Li W, et al. (2025). Post sleeve gastrectomy-enriched gut commensal *Clostridia* promotes secondary bile acid increase and weight loss. *Gut Microbes*. 17(1):2462261. doi:10.1080/19490976.2025.2462261
 - Masi D, Watanabe M, Clément K (2026). Gut microbiome and obesity care: bridging dietary, surgical, and pharmacological interventions. *Cell Rep Med*. 7(2):102573. doi:10.1016/j.xcrm.2025.102573

Ch. 28 — Cancer

- Boursi B et al. (2015). Recurrent antibiotic exposure may promote cancer formation. *Eur. J. Cancer*.
- Ugai T et al. (2022). Is early-onset cancer an emerging global epidemic? *Nat. Rev. Clin.*

Oncol. doi:10.1038/s41571-022-00672-8

- Palomar-Cros A et al. (2023). Artificial sweeteners and cancer risk. *Int. J. Cancer.* doi:10.1002/ijc.34577

Ch. 29 — Neurological and Psychiatric Disease

- Seo DO et al. (2023). ApoE and microbiota-dependent progression of tauopathy. *Science.* doi:10.1126/science.add1236
- Seo DO & Holtzman DM (2020). Gut Microbiota and Alzheimer's pathology. *J. Gerontol.* doi:10.1093/gerona/glz262
- Seo DO & Holtzman DM (2024). Alzheimer's disease-associated microbiome and therapeutic strategies. *Exp. Mol. Med.* doi:10.1038/s12276-023-01146-2
- Seo DO (2019). The microbiome: A target for Alzheimer disease? *Cell Res.* doi:10.1038/s41422-019-0227-7
- Konings B et al. (2023). GI syndromes preceding Parkinson's disease: testing Braak's hypothesis. *Gut.* doi:10.1136/gutjnl-2023-329685
- Zhu F et al. (2020). Gut microbiome and schizophrenia. *Nat. Commun.*

Ch. 30 — The Microbiome as a Diagnostic Tool

- [REF:icc2024] International consensus statement on microbiome testing in clinical practice. *Lancet Gastroenterol. Hepatol.*, 2025;10(2):154–167. doi:10.1016/S2468-1253(24)00311-X
- [REF:zeller2014] Zeller G, Tap J, Voigt AY, et al. (2014). Potential of fecal microbiota for early-stage detection of colorectal cancer. *Mol. Syst. Biol.*, 10(11), 766. doi:10.15252/msb.20145645
- [REF:wirbel2019] Wirbel J, Pyl PT, Kartal E, et al. (2019). Meta-analysis of fecal metagenomes reveals global microbial signatures that are specific for colorectal cancer. *Nature Medicine*, 25(4), 679–689. doi:10.1038/s41591-019-0406-6
- [REF:thomas2019] Thomas AM, Manghi P, Asnicar F, et al. (2019). Metagenomic analysis of colorectal cancer datasets identifies cross-cohort microbial diagnostic signatures and a link with choline degradation. *Nature Medicine*, 25(4), 667–678. doi:10.1038/s41591-019-0405-7
- [REF:pooled2025] Pooled analysis of 3,741 stool metagenomes from 18 cohorts for cross-stage and strain-level reproducible microbial biomarkers of colorectal cancer. *Nature Medicine*, 2025. doi:10.1038/s41591-025-03693-9
- [REF:zheng2024] Zheng J, Sun Q, Zhang M, et al. (2024). Noninvasive, microbiome-based diagnosis of inflammatory bowel disease. *Nature Medicine*, 30(12), 3555–3567. doi:10.1038/s41591-024-03280-4
- [REF:caussy2019] Caussy C, Tripathi A, Humphrey G, et al. (2019). A gut microbiome signature for cirrhosis due to nonalcoholic fatty liver disease. *Nature Communications*, 10(1), 1406. doi:10.1038/s41467-019-09455-9
- [REF:iebba2020] Iebba V, et al. (2020). A universal gut microbiome-derived signature predicts cirrhosis. *Cell Metabolism*, 32(5), 878–888. doi:10.1016/j.cmet.2020.06.005
- [REF:menozzi2026] Menozzi E, Schapira AHV, et al. (2026). Microbiome signature of Parkinson's disease in healthy and genetically at-risk individuals. *Nature Medicine*. doi:10.1038/s41591-026-04318-5
- [REF:oralage2026] Oral microbiome signatures predict biological age and host health. *Nature Communications*, 2026. doi:10.1038/s41467-026-72096-2

- [REF:oralpc2025] Oral microbiome as a diagnostic biomarker for pancreatic cancer: a systematic review and meta-analysis. *Acta Odontol. Scand.*, 2025. doi:10.1080/20002297.2025.2571188
- [REF:poore2020_RETRACTED] Poore GD, Kopylova E, Zhu Q, et al. (2020). Microbiome analyses of blood and tissues suggest cancer diagnostic approach. *Nature*, 579(7800), 567–574. doi:10.1038/s41586-020-2095-1 (**RETRACTED June 2024**)
- [REF:gihaw2023] Gihawi A, Ge Y, Lu J, et al. (2023). Major data analysis errors invalidate cancer microbiome findings. *Nature*. doi:10.1038/s41586-023-06593-1
- [REF:liang2024] Liang W, et al. (2024). Circulating microbiome DNA as biomarkers for early diagnosis and recurrence of lung cancer. *Cell Reports Medicine*, 5(4), 101480. doi:10.1016/S2666-3791(24)00167-8
- [REF:biomeone2023] Assessing the Performance of a Novel Stool-Based Microbiome Test That Predicts Response to First Line Immune Checkpoint Inhibitors in Multiple Cancer Types. *Diagnostics*, 2023;13(14):2329. doi:10.3390/diagnostics13142329
- [REF:blautia2025] Non-intestinal microbial signatures in stool as predictors of cancer immunotherapy outcome. *bioRxiv*, 2025. doi:10.1101/2025.05.07.652660
- [REF:park2020] Park SM, Won DD, Lee BJ, et al. (2020). A mountable toilet system for personalized health monitoring via the analysis of excreta. *Nature Biomedical Engineering*, 4, 624–635. doi:10.1038/s41551-020-0534-9
- [REF:nema2025] Nema V, et al. (2025). A small pill-like ingestible microdevice for site-specific microbiome sampling in the upper GI tract. *Small*. doi:10.1002/sml.202510289
- [REF:ebcare2025] EBCare: A smart mask for exhaled breath condensate harvesting and analysis. *Science*, 2025. doi:10.1126/science.adn6471
- [REF:bestpractices2025] Best practices for developing microbiome-based disease diagnostic classifiers through machine learning. *Gut Microbes*, 2025;17(1):2489074. doi:10.1080/19490976.2025.2489074
- [REF:mgm2024] Microbial General Model (MGM): a large-scale pretrained foundation model for microbiome analyses. *bioRxiv*, 2024. doi:10.1101/2024.12.30.630825
- [REF:microbiome_lm2023] Learning a deep language model for microbiomes: the power of large-scale unlabeled microbiome data. *PLOS Computational Biology*, 2023;19(6):e1011353. doi:10.1371/journal.pcbi.1011353
- [REF:zeevi2015] Zeevi D, Korem T, Zmora N, et al. (2015). Personalized nutrition by prediction of glycemic responses. *Cell*, 163(5), 1079–1094. doi:10.1016/j.cell.2015.11.001
- [REF:dtc_science2024] The DTC microbiome testing industry needs more regulation. *Science*, 2024. doi:10.1126/science.adk4271
- [REF:dtc_lancet2024] Direct-to-consumer microbiome testing needs regulation. *Lancet Gastroenterol. Hepatol.*, 2024;9(8):695. doi:10.1016/S2468-1253(24)00163-8
- [REF:dtc_jlb2025] Is the current regulatory framework for direct-to-consumer microbiome-based tests sufficient? *J. Law Biosciences*, 2025;12(2):lsaf024. doi:10.1093/jlb/lsaf024

Ch. 31 — Ethical, Social, and Philosophical Questions

- Spencer HG & Zuk M (2016). For Host's Sake: The Pluses of Parasite Preservation. *Trends Ecol. Evol.* doi:10.1016/j.tree.2016.02.021
- Murphy AA et al. (2016). Self-disseminating vaccines for emerging infectious diseases. *Expert Rev. Vaccines*. doi:10.1586/14760584.2016.1106942
- Rosenthal SR et al. (2015). Redefining disease emergence. *One Health*. doi:10.1016/j.onehlt.2015.08.001

Notable Observations

1. **Diabetes/metabolic disease concentration:** ~50 of the 176 references are diabetes-focused. These are concentrated in Ch. 27 (primary home) with relevant spillover into Ch. 22 (metformin), Ch. 19 (antibiotics and T1D), and Ch. 26 (autoimmune T1D). This is rich material for a strong metabolic disease chapter, but as you noted, should not disproportionately dominate the book.
2. **Strong collections for specific chapters:** Toxoplasma/behavioural modulation (Ch. 14), bacteriophages and obelisks (Ch. 16–17), the brain microbiome (Ch. 11), and artificial sweeteners (Ch. 20) all have good clusters of references.
3. **Gaps to fill:** Chapters 6 (skin), 9 (oral), 10 (mycobiome), 24 (FMT), 25 (precision engineering), and 32 (future) have few or no references in this bibliography. These will need additional literature searches during drafting.

Ch. 13 — The Gut-Brain Axis: Microbes and the Mind

- [REF:bravo2011] Bravo JA, Forsythe P, Chew MV, Escaravage E, Savignac HM, Dinan TG, Bienenstock J & Cryan JF (2011). Ingestion of *Lactobacillus* strain regulates emotional behavior and central GABA receptor expression in a mouse via the vagus nerve. *PNAS*, 108(38), 16050–16055. doi:10.1073/pnas.1102999108
- [REF:yano2015] Yano JM, Yu K, Donaldson GP, Shastri GG, Ann P, Ma L, Nagler CR, Ismailov RF, Mazmanian SK & Hsiao EY (2015). Indigenous bacteria from the gut microbiota regulate host serotonin biosynthesis. *Cell*, 161(2), 264–276. doi:10.1016/j.cell.2015.02.047
- [REF:strandwitz2019] Strandwitz P et al. (2019). GABA-modulating bacteria of the human gut microbiota. *Nat. Microbiol.*, 4(3), 396–403. doi:10.1038/s41564-018-0307-3
- [REF:sudo2004] Sudo N, Chida Y, Aiba Y et al. (2004). Postnatal microbial colonization programs the hypothalamic-pituitary-adrenal system for stress response in mice. *J. Physiol.*, 558(Pt 1), 263–275. doi:10.1113/jphysiol.2004.063388
- [REF:bercik2011] Bercik P, Denou E, Collins J et al. (2011). The intestinal microbiota affect central levels of brain-derived neurotrophic factor and behavior in mice. *Gastroenterology*, 141(2), 599–609.e3. doi:10.1053/j.gastro.2011.04.052
- [REF:diazheijtz2011] Diaz Heijtz R, Wang S, Anuar F et al. (2011). Normal gut microbiota modulates brain development and behavior. *PNAS*, 108(7), 3047–3052. doi:10.1073/pnas.1010529108
- [REF:dinan2013] Dinan TG, Stanton C & Cryan JF (2013). Psychobiotics: a novel class of psychotropic. *Biol. Psychiatry*, 74(10), 720–726. doi:10.1016/j.biopsych.2013.05.001
- [REF:kelly2016] Kelly JR, Borre Y, O'Brien C et al. (2016). Transferring the blues: depression-associated gut microbiota induces neurobehavioural changes in the rat. *J. Psychiatr. Res.*, 82, 109–118. doi:10.1016/j.jpsychires.2016.07.019
- [REF:vallescolomer2019] Valles-Colomer M, Falony G, Darzi Y et al. (2019). The neuroactive potential of the human gut microbiota in quality of life and depression. *Nat. Microbiol.*, 4(4), 623–632. doi:10.1038/s41564-018-0337-x
- [REF:asad2025] Asad A, Kirk M, Zhu S et al. (2025). Effects of prebiotics and probiotics on symptoms of depression and anxiety in clinically diagnosed samples: systematic review and meta-analysis of randomized controlled trials. *Nutr. Rev.*, 83(7), e1504–e1520. doi:10.1093/nutrit/nuae177
- [REF:hsiao2013] Hsiao EY, McBride SW, Hsien S et al. (2013). Microbiota modulate beha-

vioral and physiological abnormalities associated with neurodevelopmental disorders. *Cell*, 155(7), 1451–1463. doi:10.1016/j.cell.2013.11.024

- [REF:sharon2019] Sharon G, Cruz NJ, Kang DW et al. (2019). Human gut microbiota from autism spectrum disorder promote behavioral symptoms in mice. *Cell*, 177(6), 1600–1618.e17. doi:10.1016/j.cell.2019.05.004
- [REF:kang2017] Kang DW, Adams JB, Gregory AC et al. (2017). Microbiota Transfer Therapy alters gut ecosystem and improves gastrointestinal and autism symptoms: an open-label study. *Microbiome*, 5(1), 10. doi:10.1186/s40168-016-0225-7
- [REF:kang2019] Kang DW, Adams JB, Coleman DM et al. (2019). Long-term benefit of Microbiota Transfer Therapy on autism symptoms and gut microbiota. *Sci. Rep.*, 9, 5821. doi:10.1038/s41598-019-42183-0
- [REF:wan2024] Wan YP et al. (2024). Effect of oral faecal microbiota transplantation intervention for children with autism spectrum disorder: a randomised, double-blind, placebo-controlled trial. *Clin. Transl. Med.*, 14, e70006. doi:10.1002/ctm2.70006

Ch. 17 — Phages, Immunity, and the Tripartite Relationship

- [REF:nguyen2017] Nguyen S, Baker K, Padman BS, Patwa R, Dunstan RA, Weston TA, Schlosser K, Bailey B, Lithgow T, Lazarou M, Luque A, Rohwer F, Blumberg RS & Barr JJ (2017). Bacteriophage Transcytosis Provides a Mechanism To Cross Epithelial Cell Layers. *mBio*, 8(6), e01874-17. doi:10.1128/mBio.01874-17
- [REF:barr2013] Barr JJ, Auro R, Furlan M, Whiteson KL, Erb ML, Pogliano J, Stotland A, Wolkowicz R, Cutting AS, Doran KS, Salamon P, Youle M & Rohwer F (2013). Bacteriophage adhering to mucus provide a non-host-derived immunity. *PNAS*, 110(26), 10771–10776. doi:10.1073/pnas.1305923110
- [REF:barr2015] Barr JJ, Auro R, Sam-Soon N, Kassegne S, Peters G, Bonber N, Hatber M, Molineux S, Guardado A, Vega R, Bailey B, Pogliano J, Danthi P & Rohwer F (2015). Subdiffusive motion of bacteriophage in mucosal surfaces increases the frequency of bacterial encounters. *PNAS*, 112(44), 13675–13680. doi:10.1073/pnas.1508355112
- [REF:almeida2019] Almeida GMF, Laanto E, Asber R & Sundberg LR (2019). Bacteriophage Adherence to Mucus Mediates Preventive Protection against Pathogenic Bacteria. *mBio*, 10(6), e01984-19. doi:10.1128/mBio.01984-19
- [REF:barr2022] Barr JJ et al. (2022). Bacteriophages evolve enhanced persistence to a mucosal surface. *PNAS*, 119(27), e2116197119. doi:10.1073/pnas.2116197119
- [REF:vanbellegghem2017] Van Bellegghem JD, Clement F, Merabishvili M, Lavigne R & Vaneechoutte M (2017). Pro- and anti-inflammatory responses of peripheral blood mononuclear cells induced by *Staphylococcus aureus* and *Pseudomonas aeruginosa* phages. *Sci. Rep.*, 7, 8004. doi:10.1038/s41598-017-08336-9
- [REF:vanbellegghem2019] Van Bellegghem JD, Dąbrowska K, Vaneechoutte M, Barr JJ & Bollyky PL (2019). Interactions between Bacteriophage, Bacteria, and the Mammalian Immune System. *Viruses*, 11(1), 10. doi:10.3390/v11010010
- [REF:gogokhia2019] Gogokhia L, Buber K, Roundy J, Padgett MJ, Hapfelmeier S, Hooper LV & Round JL (2019). Expansion of Bacteriophages Is Linked to Aggravated Intestinal Inflammation and Colitis. *Cell Host Microbe*, 25(2), 285–299. doi:10.1016/j.chom.2019.01.008
- [REF:gorski2012] Górski A, Międzybrodzki R, Borysowski J et al. (2012). Phage as a Modulator of Immune Responses: Practical Implications for Phage Therapy. *Adv. Virus Res.*, 83, 41–71. doi:10.1016/B978-0-12-394438-2.00002-5
- [REF:gorski2017] Górski A, Dąbrowska K, Międzybrodzki R, Weber-Dąbrowska B, Łusiak-

Szelachowska M, Jończyk-Matysiak E & Borysowski J (2017). Phages and immunomodulation. *Future Microbiol.*, 12(10), 905–914. doi:10.2217/fmb-2017-0049

- [REF:lamybesnier2024] Lamy-Besnier Q, Theodorou I, Billaud M, Zhang H, Brot L, Culot A, Abriat G, Sokol H, De Paepe M, Petit MA & De Sordi L (2024). The human blood harbors a phageome which differs in Crohn's disease. *bioRxiv* / preprint. doi:10.1101/2024.06.04.597176
- [REF:phageometransfer2024] (2024). Phageome transfer from gut to circulation and its regulation by human immunity. *bioRxiv* / preprint. doi:10.1101/2024.06.05.597592

-
4. **Obelisks** (Zheludev et al. 2024): A fascinating recent discovery — viroid-like RNA elements in the human gut microbiome. This could warrant its own sidebar or sub-section in the phage chapters (Ch. 16–17), as it expands the concept of what biological entities inhabit our microbiome beyond bacteria, fungi, and classical viruses.
 5. **Paternal microbiome** (Argaw-Denboba et al. 2024): This is a Nature paper showing paternal gut microbiome perturbations affect offspring fitness — an important addition to Ch. 3 that expands the maternal-focused narrative of microbial inheritance.
 6. **Harada et al. (2025)**: An ultra-reduced archaeal genome — potentially relevant to Ch. 10 or a general discussion of the minimal requirements for cellular life within the microbiome.